



 Latest updates: <https://dl.acm.org/doi/10.1145/3729418>

SURVEY

Anonymization Techniques for Behavioral Biometric Data: A Survey

SIMON HANISCH, TUD Dresden University of Technology, Dresden, Sachsen, Germany

PATRICIA ARIAS-CABARCOS, Paderborn University, Paderborn, Nordrhein-Westfalen, Germany

JAVIER PARRA-ARNAU, Polytechnic University of Catalonia, Barcelona, Barcelona, Spain

THORSTEN STRUFE, Karlsruhe Institute of Technology, Karlsruhe, Baden-Württemberg, Germany

Open Access Support provided by:

Karlsruhe Institute of Technology

Paderborn University

TUD Dresden University of Technology

Polytechnic University of Catalonia



PDF Download
3729418.pdf
06 April 2026
Total Citations: 4
Total Downloads:
3830

Published: 12 June 2025
Online AM: 18 April 2025
Accepted: 02 April 2025
Revised: 19 March 2025
Received: 11 September 2023

[Citation in BibTeX format](#)

Anonymization Techniques for Behavioral Biometric Data: A Survey

SIMON HANISCH, Technische Universität Dresden, Dresden, Germany

PATRICIA ARIAS-CABARCOS, Paderborn University, Paderborn, Germany

JAVIER PARRA-ARNAU, Universitat Politècnica de Catalunya, Barcelona, Spain

THORSTEN STRUFE, Karlsruhe Institute of Technology, Karlsruhe, Germany

Our behavior—the way we talk, walk, act, or think—is unique and can be used as a biometric trait. It also correlates with sensitive attributes such as emotions and health conditions. With more and more behavior tracking techniques (e.g., fitness trackers, mixed reality) entering our everyday lives, more of our behavior is captured and processed. Hence, techniques to protect individuals’ privacy against unwanted inferences are required before such data is processed. To consolidate knowledge in this area, we are the first to systematically review suggested anonymization techniques for behavioral biometric data. We taxonomize and compare existing solutions regarding privacy goals, conceptual operation, advantages, and limitations. Our categorization allows for the comparison of anonymization techniques across different behavioral biometric traits. We review anonymization techniques for the behavioral biometric traits of voice, gait, hand motions, eye gaze, heartbeat (ECG), and brain activity (EEG). Our analysis shows that some behavioral traits (e.g., voice) have received much attention, while others (e.g., eye gaze, brain activity) are mostly neglected. We also find that the evaluation methodology of behavioral anonymization techniques can be further improved.

CCS Concepts: • **Security and privacy** → **Pseudonymity, anonymity and untraceability**;

Additional Key Words and Phrases: Privacy, behavioral data, de-identification

ACM Reference Format:

Simon Hanisch, Patricia Arias-Cabarcos, Javier Parra-Arnau, and Thorsten Strufe. 2025. Anonymization Techniques for Behavioral Biometric Data: A Survey. *ACM Comput. Surv.* 57, 11, Article 272 (June 2025), 54 pages. <https://doi.org/10.1145/3729418>

Funded by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany’s Excellence Strategy – EXC 2050/1 - Project ID 390696704 - Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of Technische Universität Dresden.

This work was performed when affiliated with KIT, funded by Helmholtz Association, topic “46.23 Engineering Secure Systems.”

This work was performed when affiliated with KIT, and funded by the Humboldt Foundation.

Authors’ Contact Information: Simon Hanisch, Technische Universität Dresden, Dresden, Germany; e-mail: simon.hanisch@tu-dresden.de; Patricia Arias-Cabarcos, Paderborn University, Paderborn, Nordrhein-Westfalen, Germany; e-mail: pac@mail.uni-paderborn.de; Javier Parra-Arnau, Universitat Politècnica de Catalunya, Barcelona, Catalunya, Spain; e-mail: javier.parra@upc.edu; Thorsten Strufe, Karlsruhe Institute of Technology, Karlsruhe, Baden-Württemberg, Germany; e-mail: thorsten.strufe@kit.edu.



This work is licensed under a [Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-nc-sa/4.0/).

© 2025 Copyright held by the owner/author(s).

ACM 0360-0300/2025/06-ART272

<https://doi.org/10.1145/3729418>

1 Introduction

The ongoing digital transformation is leading to an increasingly comprehensive data collection on citizens. Ever improving peripherals, such as **augmented reality (AR)/virtual reality (VR)** goggles, motion capturing suits and gloves, force-feedback input devices, sensor-rich cell phones, smartwatches, and other wearables drastically increase the coverage and resolution at which biometrics and behavioral data of individuals become available for processing.

A large amount of such data is shared knowingly, when users post their latest achievements, photos, or opinions on products and current affairs. A much larger amount is collected unnoticed, when individuals browse Web pages, use location services within navigation-, recommendation-, and similar apps, use wearables, or simply enter smart spaces that are enriched with anything from voice assistants to cameras.

The corresponding behavioral data is highly descriptive of the captured individual and it reveals a multitude of attributes. They contain strong indicators for routines, habits, and also medical conditions, quirks and ties. Known correlations between physiological features and medical conditions include the detection of depression [60] in facial pictures, detection of organ insufficiencies due to the coloration of eyes (hepatitis), or skin (alcohol abuse [58], general fitness [220], and others).

Behavioral data can also be used to uniquely identify individuals. Prominent examples across the spectrum include identifying personal traits and characteristics from social media feeds [145], identifying users by their mobility patterns [61], and web-browsing behavior [68]. Gait very prominently has been used to identify individuals [290, 315], and it obviously reveals individual attributes such as age, gender, and physiological conditions [282].

Preserving the privacy and ultimately the dignity of individuals who come in the range of sensors and are captured in their behavior requires more sophisticated approaches than removing direct identifiers (IP address, **Social Security number (SSN)**, blurring a face) or intuitive quasi-identifiers (gender, age, ethnicity) in databases. Note that the behavioral data captured from humans has both temporal dependencies, as it is captured as a time-series, and physiological dependencies, as human bodies must adhere to both their physiological and general physical limitations. Due to the strong dependency between observations and to the physiological and physical dependencies, the efficacy of randomized, perturbative anonymization has to critically be reviewed, as the dependencies might be used to recover the identifying information that the anonymizations seek to remove. Context information and habits being represented as strong signals in the data further complicate effective anonymization.

A growing corpus of studies is addressing this challenge of anonymizing behavioral data. They focus on a variety of different human traits, ranging from the voice, over gait, to less prominent examples such as gestures, heartbeat, and others. A systematic review of all these approaches, which bridges the attempts to extract the shared conceptual and methodological similarities, is missing, to the best of our knowledge. Further, we want to highlight the differences between approaches, their conceptual properties, as well as future research opportunities.

For this article, we hence set out to systematize the corresponding literature. We are interested in **privacy-enhancing technologies (PETs)** for scenarios in which behavioral data is collected by or shared with third parties to perform a specific operation. As we are interested more in privacy than confidentiality, we do not consider approaches in which an entity encrypts its own data to hide it from access by unintended audiences. We are rather interested in approaches that protect from unintended revelation of information contained in data [51]. We deem “confidential computing,” processing based on homomorphic cryptography, or similar approaches in which the data owner is the only entity that learns anything from the data, out of scope of our analysis. For our study, we followed Kitchenham’s guidelines [141] to systematically discover and survey the

current state-of-the-art, comprising 142 distinct studies, extracted from a corpus of 364 initially discovered publications.

We identify common applications that process behavioral data, to extract sensible measures of utility, as well as common privacy threats with corresponding adversary models. We define two taxonomies of anonymization approaches. The first is defined by how the anonymization transforms the data and the second by which anonymization goal it seeks to protect. Next, we provide a detailed overview of the different anonymization approaches, sorted by the trait they aim to protect. We provide insight into the corresponding applications that define the utility and into the privacy threats, privacy goals, applied anonymization concepts, and the evaluation the corresponding scientists performed, together with the data they chose for their studies.

As main findings, we show how the underlying anonymization concepts are independent of the biometric trait. In consequence, we identify biometric traits for which specific anonymization concepts have not yet been tested. Further, we find that the general evaluation methodology for behavioral biometric anonymization implies a weak adversary and must be improved to convincingly assess the efficacy of protective measures.

The main contributions of this work are as follows:

- Following Kitchenham’s guidelines [141], we systematically discovered a corpus of 364 proposals, which we filtered to 142 distinct proposals for the privacy protection of behavioral biometrics.
- We categorized the works by using two novel taxonomies, allowing the comparison of PETs across biometric traits.
- Further, we find that the underlying privacy protection concepts and the general evaluation methodology for behavioral biometrics are independent of biometric traits. This allows novel behavioral biometric traits to adapt concepts and methodologies from more established ones like voice.

The rest of the article is organized as follows: Section 2 describes the background on privacy terminology, as well as the related work and our survey approach. Section 3 introduces behavioral data, applications, and related privacy concerns. We define our taxonomy of concepts in Section 4, and we survey the field, sorting anonymization techniques by the trait the authors addressed and the conceptual approach taken, in Section 5. We discuss our insights and general lessons learned in Section 6 and conclude the article with a summary in Section 7.

2 Background

In this section, we first review the relevant terminology utilized throughout this work and the existing surveys on anonymization techniques. We then present the methodology we used to perform the systematic literature review.

2.1 Terminology

Our use of the term **privacy enhancement** or **protection** shall refer to the obfuscation of information from any adversarial observers, including the service provider, regardless of whether this obfuscation consists of data access control, encryption, minimization of the data revealed, or data modification, perturbation, partial or full, in any manner. In the most abstract sense, the behavioral information to be protected may be composed of various elements, including links or relationships among several pieces of information. Note that we will later focus on techniques that control disclosure in processes, where untrusted parties get access to *some* interpretable data, rather than processes in which untrusted parties get access to encrypted data only.

One important type of information to be obfuscated is a user's explicit **identity**. The close relation between personal devices (such as smartphones or wearables) and their users makes distinctive features (e.g., device fingerprints) in said devices potentially unique identifiers. We adhere to the convention that **anonymity** is the particular case of privacy in which the data may not be linked to the individual to whom the data refers to. This refers to not just direct identifiers¹ but also to indirect identifiers.

In the field of **statistical disclosure control (SDC)** [294], the aim is to protect a microdata set, while ensuring that this data is still useful for researchers. A microdata set is a database whose records contain information at the level of individual respondents. In this field, the concepts of **identity and attribute disclosure** refer to the goal of an attacker to ascertain either the identity of an individual in the microdata set or the confidential attribute/s thereof.

We shall employ the term **utility** to quantify the degree of functionality maintained concerning a service for which the behavioral biometric data is intended. The utility is kept despite the implementation of a privacy mechanism that may hide or perturb part of the data, which may degrade the quality of the service. We stress that utility in this context does not refer to user-interface design.

As pointed out above in the introduction, any PET poses a **tradeoff between privacy and functionality**. The optimization of the privacy-functionality (or privacy-utility) tradeoff will refer to the design and tuning of PETs to maximize privacy for a desired functionality or vice versa.

2.2 Related Surveys

Most of the surveys on behavioral data focus on analyzing the uniqueness and suitability of behavioral traits to identify people, comparing the accuracy of different approaches and their applicability. In this line of research, we find surveys covering a range of existing behavioral biometrics for user authentication [13, 154, 165, 182] and others focusing on the review of specific traits, such as gait recognition [290], keystrokes [19, 272], eye gaze [136], or brainwave biometrics [99]. However, the treatment of privacy issues is limited to mentioning that there is potential for sensitive inferences or identity leaks but there is no in-depth discussion about privacy countermeasures.

There is an important stream of research on potential privacy attacks to behavioral data focusing on **attribute inferences** [18, 34, 124] or dealing with user de-identification (i.e., trying to identify a person by their behavioral data) [75, 78, 108, 312]. Dantcheva et al. [53] provide an extensive overview of which sensitive attributes, so-called soft biometrics (gender, age, ethnicity, weight, etc.), can be inferred from primary biometrics extracted from image and video data. Ciriani et al. [48] performed a survey on k-anonymity, which can be used to protect from the identification of soft biometrics in tabular data. Laishram et al. [149] conducted a survey on recent advances in building privacy-preserving face recognition systems. In addition, some recent surveys focus on the privacy implications of **large language models (LLM)** [54] and generative machine learning [95, 292].

While the current literature on behavioral data underscores the need for privacy defenses, work in this area is still emerging and scattered. So far, no comprehensive view of the problem, existing solutions, and challenges has been carried out. Ribaric et al. [245] review techniques to protect user's visual and multimedia data from attribute inferences and re-identification. Though they include a section on behavioral data protection, it only covers a limited number of traits (voice, gait, and gesture) and anonymization techniques that apply when this data have been captured as video, audio, or images. No other sensors are considered. Also closely related, Nhat Tran et al.

¹Direct identifiers allow to unequivocally identify individuals. For example, it would be the case of SSNs or full names. In a data-anonymization process, direct identifiers are always removed in the very first phase.

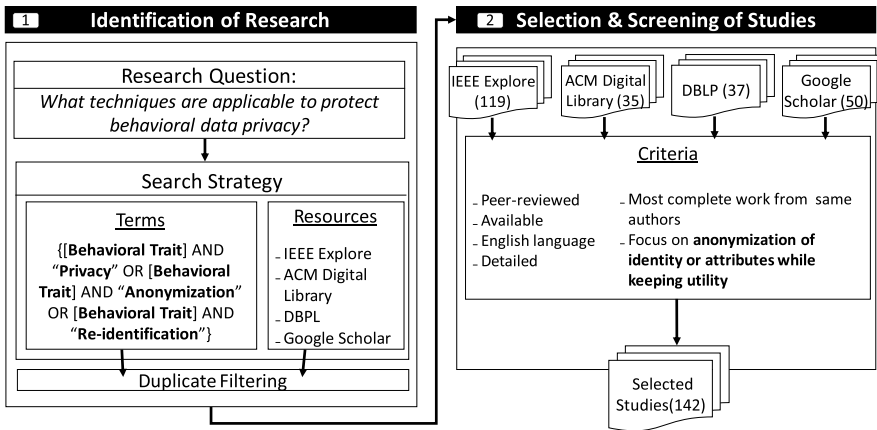


Fig. 1. Summary of the procedure for identifying and selecting relevant studies on behavioral data privacy techniques. We first analyzed the literature on biometrics to determine behavioral traits for person identification. We then used these traits as key terms to search for privacy-related publications, following Kitchenham’s guidelines for systematic literature reviews [141]. The complete list of behavioral traits we searched includes: brain activity, eye gaze, facial expression, gait, gesture, handwriting, haptic, heartbeat, keystrokes, lip, motion, mouse, thermal, touch, and voice.

[281] survey biometric template protection techniques, but they do it generally without entering in details of the anonymization needs of behavioral biometrics. Meden et al. [180] survey PETs that are applied to faces looking at different aspects such as the privacy guarantees they give and what of conceptual approaches are chosen. Shopon et al. [256] look at the wider variety of biometric traits including gait and writing style. Their taxonomy focuses on whether the anonymizations hide both the identity and attributes of the person or retain some soft biometric features.

Current reviews of behavioral biometric anonymization either consider only one specific trait or review only a few anonymization techniques. What is missing is a review that examines in depth a comprehensive set of traditional and modern types of behavioral traits for which solutions have been proposed, taking into account different types of collection sensors and use cases. In addition, a comparison of evaluation approaches across behavioral biometrics has not been done. By comparing a large set of behavioral biometrics, similarities and differences between anonymization approaches become apparent and open research questions can be identified. To address these shortcomings of the related work, we conduct a survey of anonymizations for behavioral biometric traits, compare anonymizations across traits, and also compare their privacy evaluations.

2.3 Methodology

We performed a systematic literature review following Kitchenham’s guidelines [141] to identify relevant studies on privacy techniques for behavioral data, as it is depicted in Figure 1.

Our guiding research question is **“What techniques are applicable to protect behavioral data privacy?”** From this starting point, the goal is to understand how these techniques work, what is the level of protection provided, and what are the limitations and existing open challenges. To answer these questions, we first explored the literature on biometrics [5, 13, 53, 98, 165, 182, 221, 303] to determine what kind of behavioral traits can be used to identify a person. The complete list of behavioral traits we searched includes: brain activity (also referred to as cognitive biometric), eye gaze, facial expression, gait, gesture, handwriting, haptic, heartbeat, keystrokes, lip, motion, mouse, thermal, touch, and voice. Next, we used this list of traits

combined with the keyword “**privacy**” and the semantically similar terms “**anonymization**” and “**de-identification**” as search strings in the main academic databases for computer science. Based on these search terms, we compiled works with no constraints on publication date, obtaining a set of 364 papers spanning from 2007 to October 2024 after filtering duplicates. During pre-screening, we built a taxonomy of privacy solutions and decided to narrow down the scope of the survey to anonymization techniques focused on protecting the publication of behavioral data from identity and attribute disclosure attacks. We consider approaches that assume collection, sanitization, and subsequent publishing of data, which must be anonymized but also keep a level of utility to provide behavioral data-driven services. Accordingly, the down-selection of primary studies to be analyzed in this survey considered the following criteria. Documents were excluded if:

- (1) The publication format was other than peer-reviewed academic journal or conference paper.
- (2) The paper could not be retrieved using IEEE Explore, ACM Digital Library, DBLP, or Google Scholar.
- (3) The publication language was not English.
- (4) Another paper by the same authors superseded the work, in which case the most complete work was considered.
- (5) The privacy protection technique was other than identity or attribute anonymization with data utility.
- (6) The anonymization approach was described at a high level and not enough details were provided to properly address the guiding research question.

The search and selection protocol yielded a final corpus of 142 peer-reviewed works on behavioral data anonymization, which we clustered according to the behavioral trait being protected: gait, brain activity,² heartbeat, eye gaze, voice, and hand motions (handwriting, keystrokes, mouse movements, and hand gestures). We found no papers on facial expression, lip, touch, and haptic traits that fulfill our criteria.

3 Behavioral Data Applications and Privacy Concerns

Behavioral data can be leveraged to provide valuable services for both users and companies. In this section, we summarize the application model, the main usages of behavioral data, and the related emergent privacy issues, which motivate the need for our survey.

3.1 Behavioral Biometric Data

Behavioral biometric data are a subclass of biometric data that encompasses all human behavior. While in SDC the columns of a microdata set that should be protected (e.g., name or address) are explicit, for behavioral biometrics it is not apparent which part of the data is privacy sensitive. As behavioral biometric data is captured from a human, it contains a lot of implicit dependencies between individual data points and across traits. For example, the motion of a foot is highly dependent on the motion of the corresponding leg. It may immediately imply that a person has been injured, as the behavior exhibits typical patterns of limping, although this attribute has not been made explicit in a field of the record. Another dependency to consider is the temporal dependency between data points, as behavioral biometrics are usually captured as a time-series of consequent states. These dependencies make the anonymization of behavioral biometric data challenging, as an attacker can use them to reconstruct the clear data and extract implicit disclosures from the anonymized data.

²Brainwave signals are a manifestation of both its physiological structure and the behavioral way it processes information; for example, in reaction to stimuli. In the context of this survey, we refer to EEGs as behavioral data, given that this is our main focus of study, but we acknowledge that physiological components are present.




Fig. 2. The data-publishing scenario of the survey.

3.2 Scenario

In this survey, we assume a data-publishing scenario (see Figure 2) in which the data is first transformed in a privacy protective manner and then published, or processed by, or shared with a service or application. This also includes involuntary publication, which, for example, can occur when the biometric templates of an authentication system are leaked or fitness tracker data is sold. We assume that the utility of the protected, modified data is preserved to the extent that the received service (e.g., a personalized recommendation or played virtual reality game) is still meaningful.

3.3 Applications

In general, the entire field of **human computer interaction** captures and processes behavioral biometric data, as each input over time also comprises a behavior. Keystroke patterns and mouse movement are our main input modality for computer systems today, however, new input modalities such as touch, voice, and gestures are on the rise and will likely become more relevant in the coming years. Important in this regard will be mixed reality, as it combines many of these input modalities and requires a constant monitoring of its users.

Another area where behavioral data is useful is **healthcare** and the **quantified self**. Advances in sensors and machine learning techniques enabled the development of applications for activity recognition, fall detection, and remote health monitoring that facilitate caring for elderly, sick, or disabled people and eases diagnosis [59, 212, 226]. Typical collected data is gait and motion information coming from accelerometers and gyroscopes embedded in user devices and biosignals such as heartbeat or brain activity. This data can be also processed to give health-related feedback to users, for example, to guide them through relaxation or to detect and signal cognitive states, such as being stressed, so the user can act on it.

One of the most important and well-researched application area of behavioral data is **biometric recognition** [13, 115, 165, 182]. A person's behavior, such as the way of walking or typing on a keyboard, contains unique inherent patterns that allow for verifying the identity of that person. Given that these patterns can be sensed implicitly while the person interacts with, wears, or carries a device, behavioral biometrics are generally considered more usable than other traditional biometrics like fingerprints [29, 30] and therefore a good alternative or complement to password-based authentication. Academic research has shown the feasibility of numerous behavioral traits for user authentication, to name a few: keystroke patterns [272], gait [290], touch [273], mouse movement [324], brain activity [99], or even breathing patterns [41, 42]. And some of them are already developed in commercial solutions, especially in the financial sector to prevent fraud through detecting behavior anomalies [22, 205, 283, 289].

Besides biometric recognition and healthcare, a great deal of behavioral data driven applications are focused on **personalization**. In this category, we find adaptive interfaces and services that change their content or appearance according to the predicted user preferences based on their behavior. Furthermore, personalization can be applied in many areas. To give some examples, behavioral data is used to personalize online games adapting to the player profile for a more




Fig. 3. The general behavioral-based inference process.

satisfactory experience (e.g., adjusting the level of difficulty) [328], in recommender systems to suggest online content or advertisements [241], or in education to tailor the learning experience to the student mental state (level of attention, stress, etc.) [130].

3.4 Utility

Depending on each application, the behavioral biometric data is utilized for a variety of purposes. For example, in an application for biometric authentication, an evident measure of utility is the ability to verify the identity of an individual. Likewise, in an application based on human–computer interaction, we may require the behavior to still work as reliable input modality for computer systems. In a healthcare, application, we may be interested in detecting abnormal behavior patterns and monitoring specific aspects of the behavior such as counting steps or inferring the preferences of a user for personalization. The utility of the provided service may be assessed as the performance in carrying out those tasks.

3.5 Privacy Concerns

There are also troubling privacy implications derived from the significant amount of personal information implicitly collected in behavioral data-driven applications. As we have seen, behavioral data can be used as biometrics, because it is rich in individuating information. The counterpart is that any entity that collects behavioral data could use it to identify people even if that is not the main purpose of the service they provide. What aggravates this problem is that people might not be aware that they are being measured, either because of the lack of transparency and adequate consent frameworks or because the surveillance is meant to be covert. But besides identity, behavioral data carries a wealth of potentially sensitive information that can also be abused. For example, behavioral traits such as our voice, eye gaze, gait, or brain responses are correlated with different diseases [59, 304], mental states and emotions [269, 301], and specific involuntary reactions (such as pupil dilation) can signal our interests [147].

Technically, the general process for inferring identity or other information about an individual from their behavioral data follows four steps, depicted in Figure 3. First, there is a data acquisition step in which the behavioral data is recorded and digitised. Then a feature representation that is suitable for the latter inference is extracted from the raw data. This feature representation is then usually reduced to lower the number of dimensions to reduce its complexity. In the last step, the reduced feature representation is used to perform the inference of either identity or specific attributes. Thus, machine learning techniques are applied to classify the user data as belonging to an existing user profile or not, or as belonging to a specific attribute class (man, woman). Regression models can also be applied to assign the target individual with a measure (e.g., degree of depression on a continuous scale). Based on this general workflow, a service that uses a voice-controlled personal assistant could apply the process to classify the user commanding to open an email application as the owner of the account (authentication). But it could also exploit the voice features to classify the mood of the user and offer them highly targeted advertisements, a practice that may come with discrimination and threaten user’s autonomy.

While big companies already collect a huge amount of behavioral data, the advent of affordable consumer wearables with numerous sensors (e.g., VR/AR devices with eyetracking, head pose detection, and **electroencephalography (EEG)** sensors) exacerbates the issue. Once the data is

collected, even if for a legitimate, user-consented functionality like fraud detection based on behavior anomaly, this data can be exploited to learn private information. Hence, the need for techniques to protect behavioral data is poignant. To establish a map of current research on the topic, we categorize and analyze the existing protection approaches to prevent identity and attribute disclosure.

3.6 Attacker Model

Our adversary has gained access to the behavioral biometric data of one or multiple users and now wishes to infer private information about them. The adversary has gained this access either because they are the service provider that the users have utilized, they are a user of the service and have gained the data (e.g., face images downloaded from social media), or because there has been a leak of the biometric data. As the adversary has full access to the behavioral biometric data, it can freely select an inference technique to perform privacy inferences. Further, they also might have access to additional prior knowledge about the user, such as biometric templates or soft biometrics.

4 A Taxonomy of Solutions for Behavioral Data Privacy

Based on our literature analysis, we identify two main **privacy threats** that apply to behavioral data collected/processed by a third party and can be explained in terms of the related attacker model:

- **Identity Disclosure.** The attacker’s goal is to use the behavioral data to identify the user. In this threat, we assume that the attacker is able to link the target’s behavioral data to the target’s identity and now wants to identify them in another scenario; for example, linking the user account and data in a work-related application to their account in an entertainment application. This linkage would allow the attacker to learn more about the user activity. An example of this type of attacker, as presented in Reference [265], could be a VR headset user entering a federated Metaverse offering several services (e.g., games, adult content, professional training apps). Even if the user tries to use a pseudonym when entering a foreign server, the server and other users can use transmitted behavioral data (e.g., controller/headset motions, eye-tracking) to identify the user across different pseudonyms. Moreover, it is not uncommon that behavioral data is sold to third parties or released unintentionally through a breach or hack.³
- **Attribute Disclosure.** In this threat, the attacker goal is not to re-identify the user across accounts, but to derive sensitive attributes included within the available behavioral data that the user did not intend to disclose, such as sex, medical conditions, or personal interests. The attacker might have had previous access or could have collected a dataset on which to train the machine learning model for targeted inference. For example, based on publicly available electroencephalogram datasets of alcoholic and non-alcoholic persons [134, 203], it could be possible to build a classifier that determines if newly gathered data from an entertainment application using a **brain-computer interface (BCI)** belonged to a user with an alcohol problem.

From the privacy threats, we can derive the two **anonymization goals** with which techniques can be categorized, i.e., focused on protecting user **identity** and focused on protecting specific **attributes**, as depicted in Figure 4(a).

- **Identity Protection.** The process of transforming the behavioral biometric data of a person in such a way that their identity can no longer be linked to the data. **Pseudonymization**

³<https://www.zdnet.com/article/over-60-million-records-exposed-in-wearable-fitness-tracking-data-breach-via-unsecured-database/>

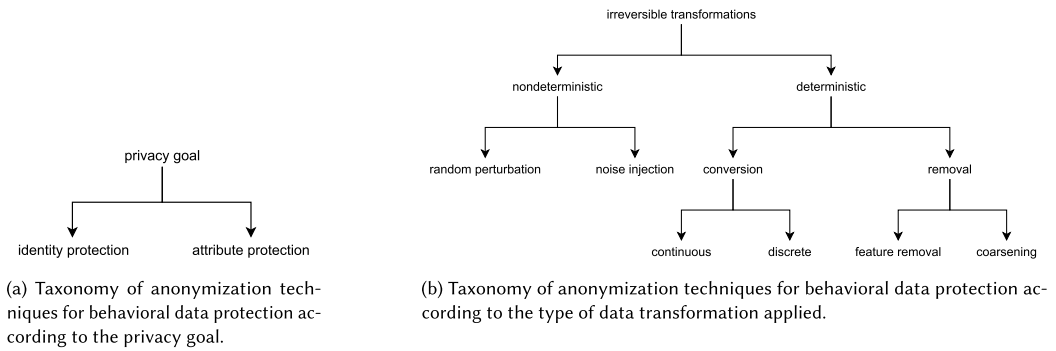


Fig. 4. Taxonomy overviews.

replaces the identifier of a person with a new one and **anonymization** prevents identification altogether.

- **Attribute Protection.** The process of transforming the behavioral biometric data of a person in such a way that specific private attributes of the person can no longer be inferred from the data. This encompasses both long-living attributes such as age or gender and short-living attributes such as mental state or temporary health conditions. An extreme version of attribute protection is template protection. For **template protection**, the identity verification of the person, in the context of an authentication system, should be still possible while all attributes are protected.

Based on the study of state-of-the-art protection methods, we have conducted a classification of methods that expectedly is not entirely exclusive to the field of behavioral data privacy, as it shares similarities with other classifications in more mature privacy fields, such as SDC. In this section, we elaborate on this classification and establish correspondences with anonymization techniques widely studied in SDC.

Our taxonomy, as depicted in Figure 4(b), of anonymization solutions for behavioral biometric data is based on the **type of transformation** applied to the original data, to derive anonymized, protected data. We include only fundamental concepts; some of the anonymization techniques combine multiple of them. The basic and shared characteristic of all anonymization methods is that they aim to provide irreversible transformations, i.e., it is impossible to transform the data back to the original data.

The first distinction of our taxonomy is if they are deterministic or randomized techniques. **Non-deterministic methods** rely on randomness in their transformation, which can yield different results for the same input, and **deterministic methods** always give the same result for the identical input. There are several methods under these two approaches, as we detail as follows:

- **Non-deterministic methods.**
 - **Random perturbations.** A random transformation into a different domain.
 - **Noise injection.** Methods that add random noise to the data points. We find that the corresponding method in the literature of SDC is referred to as *additive noise masking* [122], a perturbative technique that allows for the release of an entire microdata set, where the modified values rather than exact values are released. We would like to emphasize that additive noise masking is combined typically in this field with deterministic transformations, being it linear or non-linear.
- **Deterministic methods** are further split into **removal** and **conversion**. The removal method eliminates data points from the data such that the data points do not have an

influence on the anonymized result. Conversion methods transform the data points into a new representation, which typically depends on the original domain.

- **Removal** can be performed in two ways: **coarsening** and **feature removal**. Coarsening refers to removing parts of each data point or making the data more sparse. Feature removal refers to removing data points belonging to a specific feature altogether. This removal technique is called *suppression* [122] in the SDC field. There, when a microdata set contains too few records sharing a combination of quasi-identifier values, it is termed an “unsafe combination” due to the risk of potential re-identification. To address this concern, specific values of individual variables are deliberately suppressed and effectively replaced with missing values. This suppression strategy aims to expand the number of records that conform to each combination of key values, thereby eliminating unsafe combinations and enhancing privacy protection.
- **Conversion** can be **discrete** or **continuous**, depending on if the result of the conversion is a discrete or continuous value. As mentioned above in the noise injection technique, SDC also employs transformations of this kind.

5 Anonymization Techniques

We organize the surveyed techniques according to the behavioral biometric trait they seek to protect. We start with voice, as it stands out as the most significant trait based on available literature, and then we move on to gait, hand motions, heartbeat, eye gaze, and brain activity. For each of the traits, we analyze their utility, threat space, anonymization techniques, and evaluation methodology.

5.1 Voice

Voice processing and analysis [35] have long been performed and hence a large set of specific terminology exists to describe it. The sound of the human voice is created by the larynx and then travels via the vocal tract, which transforms and filters the sound before it leaves the mouth. Due to its approximate tube shape, the vocal tract produces resonances of the sound that are dependent on the length of the vocal tract. A phoneme is the smallest unit of sound that distinguishes one word from another, and an utterance is a unit of speech between two clear pauses. The log-spectrum is an important representation of sound, as it is closer to human perception. By using a domain transformation (**fast Fourier transform (FFT)** or cosine) on the log-spectrum, we get the cepstrum (see Figure 5). The cepstrum is useful because it allows easy estimation of the fundamental frequency (f_0) of the signal. The perceived fundamental frequency by humans is known as pitch. A widely used scale to transform the fundamental frequency to the pitch is the Mel scale. Using the Mel scale, the cepstrum can be sampled at frequencies with the same perceived distance using weighted sums. Applying an FFT on those sums gives the **Mel-frequency cepstral coefficients (MFCC)**. The MFCCs are an approximate quantification of the signal spectrum that focuses on the macrostructure of the signal.

The following gives a short overview of the field of speaker recognition (i.e., identification) that aims to establish the identity of a speaker. **Gaussian mixture models (GMM)** [243] represent speakers as the distribution of their feature vectors. The feature vectors are extracted from the speech (most often represented as MFCC) of the speaker and then modeled as Gaussian mixture density. A GMM assumes that the data points are generated by a finite number of Gaussian distributions with unknown parameters. Each feature vector is represented as a linear combination of Gaussian densities. A **universal background model (UBM)** is a GMM that models a wide variety of non-target speakers, representing possible impostors. The means of the UBM are then adjusted to the target speaker by using a maximum a posteriori adaption [244] resulting in a GMM for the




Fig. 5. A windowed speech segment (left) and its corresponding cepstrum (right), Source: <https://wiki.aalto.fi/display/ITSP/Cepstrum+and+MFCC>.

target speaker. The benefit of this approach is that the Gaussians used to model the target speaker are the same as in the UBM. For the classification of a speaker, the log-likelihood of the target speaker GMM is compared to that of the UBM to determine if the speaker should be accepted. An alternative to the log-likelihood approach is to get a GMM for each speaker recording through a **maximum a posteriori probability (MAP)** adaptation of the UBM and then map these GMM to a new feature vector called Supervector [36]. Supervectors can be classified using traditional methods like support vector machines. An extension of Supervectors is the **total variability (TV)** [64] approach. This maps the Supervectors to a low-dimensional space that models both the speaker and the channel variability. The resulting vector is called i-vector and is the de facto state-of-the-art in speaker identification. An alternative to i-vectors are x-vectors [260], which are extracted for each utterance via a **deep neural network (DNN)**.

5.1.1 Utility. The main usage of voice recordings is the transmission of information between humans, however, in recent years voice also became an important input modality for computer systems [231]. In both cases, it is important that the content of the speech is intelligible for the intended listeners. But also the mere detection of speech in audio samples can be useful, for example, for crowd detection [49]. Further, voices uniquely identify their speaker, making them suitable both for authentication and recognition purposes [247].

5.1.2 Threat Space. The privacy threats for human voices range from the identification of individuals, over the inference of private attributes, to identity theft via fake recordings. The identification of individuals via their voice has long been apparent to humans. But voices convey more information than just identity; they also allow us to infer attributes such as gender [79] or emotional state [301]. Further, modern speech synthesis methods allow the creation of fake voice recordings for a target speaker, enabling identity theft or the circumvention of speaker authentication systems. Unlike other behavioral biometric traits, voice and its resulting speech can also carry a semantic meaning, which can be sensitive to privacy.

5.1.3 Additional Privacy Goal. Voice has speech blurring as an additional privacy goal, which aims at destroying the intelligibility of the speech to protect its semantic content from unintentional listeners.

5.1.4 Anonymization Techniques. We now present the surveyed anonymization techniques that deal with protecting human voices.

Random Perturbation. Parthasarathi et al. [214] extend their feature removal methods [213] by additionally shuffling the voice blocks for adding randomness. Mtibaa et al. [193] propose a template protection scheme that relies on shuffling the feature vector of a GMM-UBM speaker identification system.

Noise Injection. Tamesue et al. [271] propose a very simple method to make speech unintelligible by simply playing pink noise between 180 and 5,630 Hz with various dBs. Ma et al. [163] also try to make speech unintelligible but focus on smartphone recordings. Their device creates two ultrasound waves whose interaction creates random low frequency waves that interfere with a smartphone's microphone but cannot be heard by humans. In their evaluation, they found that they can block smartphone recordings up to 5 meters, depending on the type of smartphone. Hashimoto et al. [107] propose a system to preserve speaker privacy in physical spaces. The core idea is to add white noise to prevent recordings of speakers from being used for identity theft. They conclude that preventing speaker identification is possible (**equal error rate (EER)** from 2% to 17%) while at the same time keeping the intelligibility of the speech at a high level (short-time objective intelligibility [270] from 1 to 0.9).

Ohshio et al. [208] train multiple so-called babble maskers from pre-recorded speakers by segmenting the speech and then averaging the segments. When a speaker should be de-identified, the babble masker is selected based on the fundamental frequency and the pitch of the person. Vaidya et al. [285] propose to add random noise to four features: pitch, tempo, pause, and MFCC. We found the descriptions of their approach to be rather short. Sharma et al. [254] use a self attention channel combinator to add noise to voice signals.

Two methods have been proposed that rely on differential privacy for noise injection. Hamm et al. [103] propose a differential private min-max filter. The min-max filter minimizes the privacy risk while maximizing utility risk with a given utility and private task. The differential privacy is achieved by adding noise either in front of the filter or after the filter. Han et al. [104] rely on X-vectors as speaker representation and formally define voice-indistinguishably as a privacy metric using differential privacy. As a measurement of similarity between x-vectors, the angular distance is used and the overall scheme gives an upper limit of this distance until which two x-vectors cannot be distinguished.

Feature Removal. Parthasarathi et al. [215] propose three feature removal methods for privacy-aware speaker change detection. Adaptive filtering assumes that the excitation source is independent of the vocal tract response. They perform short-term linear prediction analysis to estimate an all-pole model [155] (representing the vocal tract), a residual (representing the excitation source), and the gain. Then, the residual is used to estimate its real cepstrum. Their second method is to remove all subbands except the ones from 1.5 kHz to 2.5 kHz and from 3.5 kHz to 4.5 kHz. They represent the two subbands as MFCC coefficients and log-energy from a single filter. Their last method only uses the spectral slope of the speaker represented as cepstral coefficients. In another work [213], Parthasarathi et al. also propose similar feature removal methods for speaker diarisation using the real cepstrum and MFCC as features. Their analysis finds that MFCC works better than real cepstrum. Agarwal et al. [7] propose a similar scheme. They first transform the segmented speech signals into the frequency domain, then select the n most important peaks and interpolate a new signal before transforming it back into the speech domain.

Wyatt et al. [297] propose a feature removal method for speaker segmentation and conversation detection. They split the audio into segments and save for each the non-initial maximum autocorrelation peak, the total number of autocorrelation peaks, the relative spectral entropy, and the energy of the frame. Zhang et al. [322] use the same features as proposed by Wyatt et al. except for the energy of the frame and then use an HMM to perform the conversation detection. An evaluation of privacy is missing in both works.

Ditthapron et al. [69] have investigated how speech from non-target speakers can be removed in a speech assessment scenario. To separate the speakers, they first extract speaker representations from the MFCC of the speech via an encoder. The speaker representation is then concatenated

with the original MFCC before all but the target speakers are filtered out in the speaker matching network. We are missing a convincing evaluation of privacy.

Nelus et al. [201] propose to train a DNN via adversarial learning to extract features from a speaker that allow gender recognition but not speaker identification. Their evaluation shows a drop in identification from 61% to 26%, while the gender recognition only drops by 1%. They also proposed a similar system [200] that removes speaker identities from urban sound recordings. Cohen-Hadria et al. [49] also use a neural network and use it to extract the voices from recordings that consist of both background and voice noise in which the voices should be anonymized. They remove attributes with two methods. The first method simply low-pass filters the voice at 250 Hz. The second method extracts the MFCC from the voice and then uses the first five components to create a new voice. In the end, the blurred speech is recombined with the background noise. Evaluating with a speaker identification system, they were able to reduce the identification down to 29% from 43%.

Discrete Conversion. For discrete conversion, we found multiple template protection schemes.

Pathak et al. [217] present a hashing algorithm to protect voice data for authentication purposes. The supervector of a speaker is gained by performing the MAP adaptation of a universal background model for each utterance of the speaker and concatenating the means of the adapted model. The locality-sensitive hashing is then performed with the supervector that transforms it into a low dimensional space, which is referred to as a bucket. This operation is an approximation of the nearest neighbors algorithm allowing the comparison of buckets to authenticate the individual.

Portelo et al. [229, 230] propose a template protection scheme based on secure binary embeddings. The authors use a speaker identification system that uses supervectors and i-vectors to represent the features of a speaker's voice. The feature vectors are then encoded with secure binary embeddings that have the property that if the Euclidean distance of the two vectors is below a certain threshold, then the Hamming distance of the resulting hashes is proportional to the Euclidean distance. This allows the comparison of the encoded vectors by using a **support vector machine (SVM)** with a Hamming distance-based kernel.

Billeb et al. [27] propose a template protection scheme that is based on fuzzy commitment. They first extract the frequency spectrum via an FFT and then extract features from the magnitude spectrum. Then, the MAP adaptation of a GMM-UBM speaker identification system is applied and additional statistics are extracted. The template is then stored as a combination of error-correcting code and hash algorithm.

Continuous Conversion. Most voice anonymization techniques fall into the category of continuous conversion, since they attempt to create an anonymized speech recording. We have found the following techniques:

Speaker transformation is the process of manipulating the voice characteristics of a speaker (not the linguistic features) to make the voice sound like a target speaker. A target speaker can be either a specific natural speaker or a synthetic speaker. For the synthetic speaker, either an existing speaker is used or a new one is generated, for example, by averaging multiple speakers into one. The general approach of speaker transformation is that the voice characteristics of the source speaker are extracted and then transformed to match the target speaker. In the last step, the new speaker is synthesized. The following methods perform speaker transformation:

Jin et al. [129] evaluate four methods for speaker transformation for identity protection. Their base method uses a GMM-mapping-based speaker transformation system to transfer speakers to a target synthetic voice called kal-diphone. Further, they test duration transformation in which the length of utterances of the source speaker is scaled to match the ones of the target speaker. Last,

they try an extrapolated transformation in which they use the linear mapping of the source to the target to extrapolate beyond the target. Pobar et al. [225] also use a speaker transformation system based on GMM mapping but combine it with a harmonic stochastic model. Instead of retraining the system for a new speaker, one of the existing transformation functions is applied. This removes the need for a parallel corpus for the speakers that should be protected. The target speaker is a synthetic speaker that reduces the identification accuracy from 97% down to 9%.

Justin et al. [132] investigate the intelligibility of transformed speakers. They test with a diphone speech synthesis system and an HMM-based speech synthesis system to transform speakers into a synthetic speaker. They performed a survey with human listeners to evaluate the intelligibility of the protected speakers, measuring the word error rate. Abou-Zleikha et al. [3] do not propose a speaker transformation method themselves but explore how to select a target speaker to achieve the lowest identification rate and have good results when the speaker is transformed back to the source speaker. They formulate this as an optimization problem and measure the distance between two speakers with a confusion factor, for which they evaluate entropy and Gini index as metrics. Pribil et al. [234] propose a speaker de-identification method that relies on modifying several features of the source speaker. In the first step, the prosodic and spectral features are extracted from the source speaker. They then modify the features to make the speaker sound older, younger, more female, and more male.

Bahamanienezhad et al. [17] have developed a speaker transformation method that uses a convolutional encoder/decoder network. They first extract spectral features and excitation features (f_0) from the source speaker. The spectral features are then mapped via the encoder/decoder framework to a target speaker. The resulting speech is fused either via taking the average or via a gender-based average to create an average speaker. From the excitation features, only the fundamental frequency is transformed via linear transformation; the remaining features stay the same.

Fang et al. [82] use a similar averaging approach but rely on x-vectors. They extract the x-vector of a speaker and then use a set of random x-vectors of unrelated speakers to calculate a mean x-vector. In their evaluation, they demonstrate EER up to 34% for their anonymization. Mawalim et al. [177] propose to improve the system by Fang et al. by scaling the f_0 frequency either up or down, increasing the length of the speech utterances by 1.2, and using singular value modification for the combination of the x-vectors. Their EER improved up to 54%. Further improved was this system by Prajapati et al. [232], who added a CycleGAN to modify the speakers. Cheng et al. [45] propose another speaker transformation that uses one encoder for content and one for the speaker identity and then recombines them in a single decoder into the anonymized utterance. Panarielle et al. [210] use **neural audio codecs (NAC)** for the speaker transformation. Similar to other speaker transformation techniques, they independently encode the content and the speaker identity and then combine them using transformer models before decoding them using the NAC decoder.

As more speaker transformations appear, some of them focus on specific subproblems of speaker transformation. Miao et al. [186] developed a speaker transformation that is language-independent. The architecture of the system is based on the B1 baseline of the VoicePrivacy [279] challenge and they are able to show that their system works on both English and Mandarin speaker datasets. Hintz et al. [110] investigate how to anonymize stuttering speakers using a GAN to preserve the pathology of the stuttering intact while removing the speaker identity. Yang et al. [306] have developed a low-latency speaker transformation technique. Yao et al. [308] attempt to improve the distinctiveness of anonymized speakers by scaling the formant and pitch information. Meyer et al. [185] propose a speaker transformation method that preserves the prosody of the speaker. Nespoli et al. [202] propose to use two speaker transformation systems in a row to achieve better anonymization results.

Several papers investigate how the target speaker for speaker transformation can be either selected or created using the original speaker as a starting point. Chang et al. [40] and Meyer et al. [184] investigate different averaging strategies. Yuan et al. [316] train an autoencoder and use it for synthetic data generation to generate random speakers. Lv et al. [162] use autoencoders to obtain a latent representation of the speaker and then select similar latent representations from a pool using k-means. Yao et al. [307] encode the speaker as a matrix. This matrix is then decomposed using **singular value decomposition (SVD)** into eigenvectors and a matrix that stores the importance of each eigenvector. They then use a logarithmic transformation to make the importance values more similar before reconstructing the speaker identity matrix. Miao et al. [187] extend their method [186] by removing the speaker pool and using an adversarial perturbation to transform the speaker vector. Perero-Codosera et al. [219] also propose an approach using an adversarial perturbation for anonymizing the original speaker X-vector, and Yao et al. [309] propose removing random dimensions of an X-vector to create a new speaker identity.

Adversarial Perturbation: In recent years, the technique of adversarial perturbation has become popular. The general idea is that the anonymization is performed by a machine learning system that is trained with two losses. One loss is for the privacy attribute to be protected and should be minimized, while the other loss is for the desired utility and should be maximized.

Cheng et al. [43] propose VoiceCloak, which trains a convolutional perturbation injector to take the room impulse response and the original voice signal as input and outputs an anonymized voice. Deng et al. [65] present V-Cloak, which uses a convolutional autoencoder trained to minimize identification while preserving the timbre and intelligibility of the utterance high. Unique to the system is that it feeds the down-sampled speech into most of the layers of the autoencoder. In their evaluation, they can show that their system can be used for real-time anonymization and performs better or as good as other state-of-the-art voice anonymizations.

Chouchane et al. [47] address fairness concerns in speaker verification. They use adversarial training to create a speaker verification system that produces speaker embeddings that can still be used for speaker verification but no longer work for sex recognition. Xiao et al. [299] have developed a microphone module that anonymizes the speaker by adding an adversarial perturbation to the sound signal encoded by a generic **code excitation linear prediction (CELP)** codec. An interesting distinction from other adversarial perturbations is that they use a genetic algorithm to find the adversarial perturbation rather than gradient descent. They also show that their approach adds very little latency overhead. Ravi et al. [240] developed an adversarial perturbation for the utility goal of depression detection in speakers.

Ali et al. [9] also propose an autoencoder to anonymize at the network edge specifically for the input of voice assistants. Their idea is to extract privacy-friendly features by training classifiers on the latent code of the voice samples. They use the trained classifier to perform gradient reversal on the encoder to unlearn the features learned for identity, gender, and language. Yoo et al. [314] use a CycleGAN for speaker anonymization that uses a variational autoencoder as its generator. They train against a DNN speaker recognition system as the discriminator.

Frequency warping is a technique that is similar to speaker transformation; the main difference is that frequency warping focuses on transforming the frequency spectrum of a speaker and usually does not try to transform the source into a specific target speaker. It is mostly used for identity and gender protection. A common goal of frequency warping is vocal tract length normalization in which the resonances that are specific to an individual's vocal tract length should be removed or altered.

Faundez-Zanuy et al. [83] explore two approaches for gender protection: Phase vocoder and vocal tract length normalization. The vocoder approach detects peaks in the voice signal. For each

peak, a bin is defined and compared to its two neighbors to define a region of influence. Then, the peak and its region of influence are shifted by a peak specific frequency. For both genders, they can reduce gender recognition to chance level, however, the identity recognition is also close to chance level. Valdivielso et al. [1] present a speaker protection approach that transforms the pitch and the frequency axis. Lopez-Otero et al. [161] rely on frequency warping and amplitude scaling for speaker protection in the context of depression detection. They implement both operations as an affine transformation in the cepstral domain and manually define piece-wise linear transformation functions. They demonstrate an increase of the EER from 9.7% to up to 44% for the speaker identification, while the depression detection stays similar to the clear data.

Magarinos et al. [164] also rely on frequency and amplitude warping for speaker protection. First, they extract the cepstral voice vectors from the speaker and then convert them into a discrete spectrum. Then, **dynamic frequency warping (DFW)** is applied to map the source spectrum bins to the target spectrum. As multiple source bins can have the same target bin, all source bins that map to the same target bin are averaged. Additionally to the frequency and amplitude warping, the fundamental frequency is adjusted regarding its mean and variance. They demonstrate an identification reduction from 99% to 4%.

Aloufi et al. [11] try to hide the emotional state of speakers before their speech is sent to a voice-based cloud service. They first extract the fundamental frequency, spectral envelope, and aperiodicity. The features are then transformed via a CycleGAN from emotional speech to neutral speech. Their results for hiding the emotional state show a reduction from over 70% to about 20% and for hiding sex a reduction from up to 99% to the chance level of 50%.

Srivastava et al. [263] evaluate multiple speaker protection methods against an informed attacker. They work with three attacker models: an ignorant attack that is not aware that the voice data is de-identified, a semi-informed attacker that knows that the data is de-identified, and an informed attacker that knows the de-identification method and its parameters. The first method is a vocal tract length normalization approach. The speaker is represented as a set of centroid spectra. The algorithm then calculates the closest path between the source set and the target set to get the parameters for the warping. The second method uses a neural net encoder/decoder approach to transform the speaker. They found large differences for the different attacker models: While the ignorant attacker can achieve EER of up to 50%, the informed attacker only achieves 11% as its highest EER. This finding highlights how important strong attacker models are for the evaluation of anonymization techniques.

Patino et al. [218] pseudonymize speakers by transforming their McAdam coefficients. In the first step **linear predictive coding (LPC)** is applied to an input speech frame. The coefficients of the LPC are then transformed into poles, and the poles that have a nonzero imaginary part are shifted according to the angle between the real and imaginary part of the pole. Their evaluation shows that this approach performs well against an ignorant attack that is not aware of the anonymization increasing ERR from 3% to 26% while an informed attacker still achieves 5% ERR. Gupta et al. [100] further improve on transforming the McAdams coefficients by not only changing the angle of the complex poles but also modifying their radius.

Mawalim et al. [178] propose two frequency modifications for voice anonymization. Their first technique segments the speech signal and then resamples the segments to raise or lower the pitch. They then use a Hann window function to combine the segments into the speech signal. Their second technique uses a different recombination technique by recombining the overlapping segments using phase propagation. Gaznepoglu et al. [92] modify the B1 baseline of the VoicePrivacy challenge [279] to produce better anonymized fundamental frequencies by first extracting them from X-vectors and then using a mask to anonymize them.

Continuous Conversion + Random Perturbation. Canuto et al. [39] propose a new method for template protection in which the feature vector is shuffled via a randomized sum. For each feature vector, the elements are shuffled based on a secret key. Two random vectors of the same length are derived from the key. These vectors give the position of the attributes that should be summed. The reorganized feature vector is summed up with the vectors resulting when the position vectors are applied to the original feature vector.

Prajapati et al. [233, 257] first use a regular voice conversion system and then perturb the speed of the speech sequence by changing the length of the sequence. They also adjust the tempo of the sequence by cutting the sequence into segments and making them randomly shorter or longer. They recombine the segments by using a overlap-add method. Their evaluation shows that the speed perturbation makes the anonymization stronger.

Continuous Conversion + Noise Injection. Kondo et al. [143, 144] create so-called babble maskers by segmenting speech into 10-second segments and then averaging them into babble maskers. Besides speaker-dependent maskers, they also create gender-based babble maskers based on multiple speakers of the same gender. The babble masker is then applied to the recording of the speaker. Qian et al. [236] present a method to sanitize speech before it is sent to the server of a virtual assistant. Their main method is to perform vocal tract length normalization via a compound frequency warping function consisting of a bilinear and a quadratic function to avoid re-identification attacks. Additionally, they add Laplace noise after the warping function to make the anonymization more robust. For the result, they claim to achieve differential privacy. In a follow-up work [235], the same authors further investigate the security of their scheme. Srivastava et al. [263] also investigate the security of the scheme with stronger attackers.

Shmsabadi et al. [253] aim to provide theoretical privacy guarantees for speaker transformation. They do this by adding differential privacy to the pitch and context features used in speaker transformation. Both features are encoded by a specific autoencoder network, which transfers them into their latent space. For the pitch, they then add Laplace noise and then perform a clipping of the latent vector values before decoding back to pitch space. For the context features, they first normalize the latent vector and then add the Laplace noise before normalizing again and then decoding back. Due to the correlations between speech segments, it is unclear whether the differential privacy guarantees hold.

5.1.5 Evaluations. Most of the reviewed works evaluate the quality of the de-identification by comparing the recognition rates of attributes or identities on unmodified and de-identified data. The recognition is done via machine learning models or human listeners. As metrics to measure the recognition rate, the papers mostly rely on the **equal error rate (EER)**, **false positive rate (FPR)**, **false negative rate (FNR)**, recall, precision, and F1 score. Abou-Zleikha et al. [3] also use entropy and the Gini index to evaluate the de-identification performance. We believe that the prevalence of EER shows that the underlying scenario focuses on speaker verification scenarios, but we believe that speaker identification is a more appropriate scenario for evaluating speaker anonymization.

Additionally to the de-identification, some works evaluate the loss of utility. One important goal in regard to human listeners is to achieve a natural-sounding de-identified voice. The naturalness is evaluated by human listeners using the mean opinion score. Another important aspect is the intelligibility of the de-identified speech. Intelligibility can be evaluated via human listeners or machine learning models using the word error rate, phoneme error rate, or short-time objective intelligibility. A common limitation we observed is that most evaluations use the clear data to train the recognition model and then test it against the anonymized data. This approach implicitly assumes that the attacker is not aware of the anonymization and hence does not try to circumvent it.




Fig. 6. The phases of the gait cycle. Source: Reference [267].

It is worth noting the VoicePrivacy challenge [279], an initiative to improve the methodology of speaker anonymization. They use EER and the **log-likelihood-ratio cost function (Cllr)** to evaluate speaker verifiability and word error rate to evaluate speech intelligibility. In a post evaluation, they also retrained their speaker verification systems with anonymized speech data to test against an informed attack. In recent years (since 2020), the VoicePrivacy Challenge framework has become a popular choice for evaluating voice anonymization. The baselines of the challenge have also often been used as the basis for new anonymization techniques.

Qian et al. [237] present a framework to reason about the privacy and utility of voice anonymization techniques. They present the measure of p-leak limit that should give a maximum privacy leakage per speaker for a published dataset. Zhang et al. [321] propose a theoretical framework to quantify the privacy leakage risk and utility loss for speech data publishing.

5.2 Gait

The human gait is the pattern in which humans move their limbs during locomotion. Multiple manners of gait exist, such as trotting, walking, or running. Gait can be broken down into individual gait cycles [267] (see Figure 6), which is the shortest repetitive task during the gait. The gait cycle spans from a specific gait event of one foot until the same foot reaches the same gait event. It consists of a stance phase, in which the foot is on the ground, and a swing phase, in which the foot is in the air. The two phases alternate for each foot. Due to its usefulness as a behavioral biometric trait for identifying individuals, gait has long been a research interest of both computer science and psychology. For example, Yovel et al. [315] find that it plays an important part for humans to identify people at a distance, and Pollick et al. [227] show that it is possible for humans to infer the gender of a walker, even when the walker is only shown as a set of points, as so-called point-light-display. The following section deals with the anonymization of gait patterns.

Gait recognition methods have been an active research topic in the past, hence a large set of different methods for various capture methods exists. Wan et al. [290] performed a recent survey on the subject and listed recognition methods for cameras, accelerometers, floor sensors, and radars. The main portion of the works focuses on camera-based gait recognition, which is classified by Wan et al. as either model-based or model-free. Model-based methods use a specific model of the walker, for example, a pendulum model of the legs, to then match the walker to it. Model-free methods, however, do not have an explicit model but rather use the entire capture of the gait to perform the recognition, for example, by averaging the silhouette of the walker over time as a gait energy image. Accelerometer-based systems also average the gait into a feature representation either by segmenting the gait into its gait cycles or by using frames with a fixed size.

5.2.1 Utility. Gait recordings are important for medical diagnosis of gait abnormalities [140]. Another more casual example would be the recording of the gait pattern to count the steps a person has performed during a day [261]. Further, gait patterns are often recorded in videos; to not degrade the quality of the video, the gait should appear natural and convincing to its viewers [125].

5.2.2 Threat Space. Due to its omnipresence in everyday life, human gait is easy to capture, especially because most capturing methods are unintrusive and do not require the participation of the victim. Additionally, it has been shown that gait recognition is very robust to video quality and obfuscation, making it very much suited for surveillance systems [290]. Besides identifying humans, it has also been shown that gait can be used to infer private attributes like gender [227]. Considering all this, the threat to gait biometrics is already large. What is more, with recent developments in richer capturing methods such as LiDAR [87] or cheap motion capture suits, it is to be expected that the threat space for gait will even increase in the coming years.

5.2.3 Anonymization Techniques. In the following, we present the gait anonymization methods found in the literature, sorted by our taxonomy.

Random Perturbation. Hoang et al. [113] propose a fuzzy commitment scheme based on **Bose-Chaudhuri-Hocquenghem (BCH)** codes for storing accelerometer gait templates. After the feature extraction and binarization of the accelerometer data, the reliable bits are extracted. These bits are then XORed with the BCH encoded secret key to gain the secure γ . Additionally to the γ , the hash of the secret key and some helper data are stored. During the authentication phase, the extracted reliable bits are XORed with the secure γ and then decoded with BCH. The result can then be hashed and compared to the hash of the secret key. While the false accept rate is promising, the false reject rate of this scheme must be improved to be more user-friendly.

Noise Injection. The influence of noise injection on the performance of accelerometer/gyroscope authentication systems was studied by Matovu et al. [175]. For their approach, they generate a time series of noise values drawn from a uniform distribution and then merge the original time series with the generated one.

A noise injection approach for gait in videos was developed by Tieu et al. [276]. They use a **convolutional neural network (CNN)** to mix the gait of a second person (noise gait) into the original gait. In the first step, the silhouette for both the original and noise gait is extracted from a black-and-white representation of the input videos. The noise gait is selected hereby to have the same size and view angle as the original gait to achieve a more natural result. The silhouettes are then fed into the CNN, which uses shared weights networks to abstract them and then merges the abstracted representations via a third network. In a post-processing step, the original gait is replaced with the newly merged gait. Depending on the view angle, they achieve identification rates between 20% and 1%.

The authors further improve their method in a follow-up paper [277]. Here, the noise gait is generated via a **generative adversarial network (GAN)** that takes Gaussian noise as input and outputs noise silhouette. Instead of using a CNN, they then use a **self-growing and pruning GAN (SP-GAN)** to fuse the noise and original gait. Here, the identification accuracy was between 30% and 10%. Further, they propose an approach to colorize the resulting black-and-white silhouette [278]. Hanisch et al. [105] investigated multiple anonymization techniques to protect identity and gender of walkers recorded via motion capture suits. One of their techniques was to add Laplace noise to all body positions of the walker, however, their results show that effectively anonymizing was not possible without destroying the utility (measured as naturalness via a user study). Another paper that performs simple noise injection is by Meng et al. [183]; they also show that the noise level required for effective anonymization destroys the utility of the data.

Coarsening. Nair et al. [198] experiment with coarsening the frame rate, positional accuracy, and dimensionality of VR motion data. They find that while these techniques can reduce identification rates for individual motion sequences, they do not allow effective anonymization on a per-session basis and are therefore not effective for anonymizing motion data.

Feature Removal. A feature removal approach for privacy-preserving activity recognition via accelerometers is proposed by Jourdan et al. [131]. They extract various temporal and frequency features from the accelerometer data such as mean, correlation, energy, or entropy. Via experiments, they then determine the influence of each feature for activity and identity recognition. They find that the temporal features contribute more to identity recognition and the frequency features contribute more to activity recognition, therefore they remove the temporal features. Their results show a good tradeoff between activity recognition (96% reduced to 87%) and identification (90% reduced to 40%).

Debs et al. [63] do a similar simple feature removal approach, but they first transform the signal using a short-time Fourier transformation before randomly removing 10% to 90% of the data. Garofalo et al. [90] propose a temporal convolutional network as feature extractor that is trained via adversarial training. After the feature extractor created a feature vector, it is evaluated by an identity verifier and an attribute classifier, which results are then used as the loss function for the feature extractor training. Rouge et al. [246] developed an anonymization technique for accelerometer motion data. Their technique is to first extract appropriate features from the raw data using a short-time Fourier transform. They then train a random forest classifier to perform action and identity recognition. Using the trained random forest model, they then determine the importance of the features for both classification tasks and remove those that are only important for identification.

Another technique tested by Hanisch et al. [105] was to remove body parts from gait motion capture data to see their impact on the recognition of identity and gender. They found that the gait data is very redundant, and even when only the data for the head is kept, identification success remains close to 60%.

Continuous Conversion. A continuous conversion approach is blurring, in which persons in videos, including their gait, should be de-identified. As a first step, the silhouettes of the persons in the videos are tracked and segmented to then apply the blur. Agrawal et al. [8] proposed two blurring approaches: exponential blur and **line integral convolution (LIC)**. Exponential blur regards the video as a 3D space with the time as the z-axis and then calculates a weighted average of the neighbors of each voxel to blur via an exponential function. LIC works with the bounding box of the walker silhouette and maps it onto a vector field that is then used to calculate the output pixels. Another blurring approach is proposed by Ivacic-Kos et al. [125]. They apply a Gaussian filter to blur the silhouettes of walkers. The filter calculates a weighted average of the color of the neighboring pixels, with the weights decreasing monotonically from the central pixel.

Halder et al. [102] work on gait anonymization in videos. They first extract the gait silhouettes from a large number of videos. They then perform a k-means approach to cluster the silhouettes to generate a database of key gait poses. To anonymize a given gait sequence, they also extract the gait silhouette and match it to the closest key pose in the database. The key pose sequence is then used to generate a new video sequence using a GAN.

Moon et al. [191] investigate the use of adversarial training for anonymizing motion data. They train different machine learning models on 3D pose data to maximize action recognition while minimizing the identification. Their evaluation on the ETRI-activity [127] and NTU60 [157] datasets shows that they can achieve both a high utility for the action recognition and identification rates close to chance. Nair et al. [196] also propose an adversarial approach for the anonymization of VR motion data using a Siamese architecture for the training of the anonymization. Instead of using only the motion sequence as input, they also add a random vector. As before, they train their model to achieve good action recognition and low identification.

Thapar et al. [275] consider the anonymization of gait in egocentric videos, which are videos that are recorded from a first-person perspective. They first learn the identities of gallery videos

via the rotation of the camera, which is then transformed into the camera rotation signature via guided backpropagation. This camera signature is then applied to the target video, mixing the gallery identity and the target identity. In their evaluation, they test the identification of persons and find that the EER increases from around 20% to around 50%.

Continuous Conversion + Discrete Conversion. An approach that combines both continuous and discrete conversions for walkers in videos is proposed by Hirose et al. [112]. First, they extract the silhouette and the gait cycle of the walker. The silhouette is then transformed via a deconvolutional neural network encoder into a silhouette code. The code is converted by using a k-same approach in which the k-nearest neighbors of the input code are selected and then a weighted average is computed. The gait cycle is transformed via a continuous, differentiable, and monotonically increasing function. In the last step, the new video is generated by feeding the perturbed silhouette code and gait cycle into the convolutional neural network decoder. Their evolution shows that the gait recognition drops from about 100% down to 29%, 21%, and 4%, depending on the recognition model.

5.2.4 Evaluation. Gait de-identification is evaluated in the literature via gait recognition systems or human observers with the recognition accuracy as the main metric, but there are also usages of the F1 score, **equal error rate (EER)**, or **false acceptance rate (FAR)**. To access the utility loss, there is a larger variety of metrics, usually to either quantify the naturalness of the de-identified gait or to perform another kind of recognition, such as activity. One specific evaluation method we observed was by Matovu et al. [175], in which the authors used the biometric menagerie to observe the de-identification influence on different types of users in biometric authentication systems.

5.3 Hand Motions and Gestures

We use the term hand motions as an umbrella for all hand-motion-related biometric factors, including handwriting, keystrokes, mouse movements, and hand gestures. These traits mostly differ by how they are recorded and what kind of hand motions are performed. Handwriting can be captured offline or online, depending on if only the resulting written text or a real-time capturing of the hand while writing is being used. For this survey, we only consider the uniqueness of one writing style and not the linguistic style (Stylometry) of the written text. In modern life, handwriting has been mostly replaced by typing on keyboards, which also is an important biometric factor, as individuals can be identified by the timings of their key presses. Besides keyboards, also the usage of computer mice creates unique patterns, as their trajectories and clicks are again a biometric factor. Last, hand motions can be directly captured using optical or accelerometer tracking techniques.

Hand motion recognition encompasses multiple recognition techniques for different capture modalities. Here, we give an overview of handwriting, mouse movements, keystrokes, and gestures. For handwriting based hand motion recognition, the input handwriting sequence is often adjusted for its baseline, scaled to a normal writing style, and segmented to meet the demands of the classifier [223]. Handwriting is further dependent if it was captured while the person was writing (online handwriting), for example, with a digital pen, or only handwriting itself is captured after the person has finished (offline handwriting). The recognition for mouse movements relies on the trajectory, speed, single, and double clicks performed with a mouse as features. Keystroke-based hand motion recognition is based primarily on the timing differences between key up, down, and hold events. Besides individual events, the differences between two successive events or even three successive events are also used as features [327]. Hand motion recognition via gestures can be split into 2D gestures that are performed on a flat surface (e.g., on a smartphone) and 3D gestures that are performed

in mid-air. Sherman et al. [255] use the trajectories of each finger and first resample them using a cubic spline interpolation to get a lower sampling rate, removing unwanted jitter. To calculate the distance between two gestures, dynamic time warping is employed with various distance metrics.

5.3.1 Utility. The utility range for hand motions is large and diverse. For handwriting, the resulting text must be readable either by humans or computers; the particular handwriting style is usually not important. This is different for signatures, as their main purpose is to facilitate the identification and verification of the signer's identity, hence their particular style is important, while the readability of the name is less important. Since the other hand motions mostly serve as input modalities for computer systems, their utility as input modality [320] must be kept precise and timely to keep their utility. For hand gestures [250], there is additionally its utility for non-verbal communication.

5.3.2 Threat Space. The threat space for hand motion is diverse, as the usage of our hands is unavoidable in most everyday tasks, and as we often use digital devices, the recording of hand motions happens most of the time without us realizing it. As many studies have shown, hand motions can be used to identify individuals by their handwriting [223], keystroke dynamics [12], mouse movements [242], and gestures [305]. Besides identification, our hand motions also often convey meaning, such as when we type text on a keyboard; the semantics of hand motions can also be sensitive too, such as when we enter passwords or write private messages. Specific medical conditions manifest themselves in hand motions, such as hand tremors in Parkinson's patients [128]. Further, hand motions convey information about our emotional state [264].

5.3.3 Anonymization Techniques. In the following, we present the suitable methods for hand motion anonymization, with the exception of mouse movements, as we did not find any suitable papers for it.

Random Perturbation. Maiorana et al. [167] propose a template protection method for online handwriting that splits a handwriting sequence into segments and then randomly mixes the segments before convoluting them. The same shuffling approach is taken by Maiti et al. [168] to prevent keystroke inference attacks via wrist-worn accelerometers, however, they do not convolute the segments. The approach was only evaluated with four participants. Another study investigating the permutation of keystrokes is performed by Vassallo et al. [288]; in their evaluation, they only investigate the utility reduction. Goubaru et al. [97] propose a template protection scheme for online handwriting templates. They extract the pattern ID for a user by using a common template. The pattern ID is then XORed with a secret that was encoded by an error-correcting code. The result is stored as the template. For the verification, the pattern ID is again extracted and then XORed with the template.

Noise Injection. Migdal et al. [188] add delays to keystroke timings. Shahid et al. [252] propose to use the Laplace mechanism on the 2D coordinates of handwritten text to achieve local differential privacy.

Coarsening. Vassallo et al. [288] explore suppression of keystrokes to preserve the content of the typed text in a continuous authentication scenario. Maiti et al. [168] also focus on keystrokes privacy and propose two coarsening methods to prevent keystroke inference attacks via wrist-worn accelerometers. In their first approach, they simply detect if a user is typing via several features and then block the access to the accelerometer data to prevent attacks. Their second method reduces the sampling rate of the accelerometer.

Discrete Conversion. For discrete conversion, we found the following techniques aimed at template protection: An online handwriting template protection scheme is proposed by Sae-Bae et al. [248] that decomposes signatures into histograms on which the authentication is performed. They use one-dimensional histograms to capture the distribution of single features and two-dimensional histograms to capture the dependence between two features. Migdal et al. [189] propose a template protection scheme for multiple modalities, including keystrokes. Their scheme combines multiple pieces of information, such as IP addresses, with the keystroke information and then computes a bihash on it. Leinonen et al. [151] investigate the anonymization of keystroke timing data using two rounding approaches that effectively sort the timings into buckets. Their approach appears to be effective, as the identification drops from close to 100% to below 10%. Vassallo et al. [288] explore substitution of keys with a random nearby key to preserve the content of the typed text in a continuous authentication scenario.

Figueiredo et al. [85] have developed a modeling language that can be used to design new gestures for applications. The gestures can then be recognized on the recording hardware, eliminating the need to give the application access to the clear data. No privacy evaluation was performed. For privacy-friendly gesture recognition, Mukojima et al. [194] designed a system that illuminates the hand with a random pixel pattern and captures the remaining light on the opposite side of the hand with a detector. From this reduced data collection, the shape of the hand is reconstructed via machine learning. The authors did not evaluate the privacy protection of their approach.

Continuous Conversion. Maiorana et al. [167] propose two continuous conversions for online handwriting templates: a baseline conversion that first splits a handwriting sequence into multiple segments based on a secret key and then convolutes the segments; and a shifting transformation that applies a shift to the initial sequence. The template matching is performed on the protected template. For the anonymization of gestures that have been captured via **inertia measurement unit (IMU)** sensors, Malekzadeh et al. [171] propose two separate autoencoders. The first autoencoder is supposed to replace sequences in the data that have been classified as sensitive with a generated neutral sequence, while the second one should minimize the mutual information between the data and the identity of the user. Their approach reduces the identification from 96% accuracy down to 7%. Fan et al. [81] also propose using two encoders. They use one for task encoding and one for identity encoding and then feed both encodings into the decoder. This system is trained in an adversarial approach to reduce identity recognition and increase action recognition using a small sEMG dataset.

Another auto-encoder-based approach is proposed by Saunder et al. [250], in which the sign language motions of one person are transferred onto another one. Their technique is two-fold. They first extract the pose of the source video and encode this to a set of pose features. Second, they encode the style of the target appearance using an appearance distribution. The encoded pose and style are then combined to generate a new image. It was not evaluated if the persons can be identified by their hand motions only. A second approach to perform sign language anonymization was proposed by Xia et al. [298]. They use an estimation of the motion regions and then use optical flow in combination with a confidence map to encode the motions of the source and driving video. Then, the anonymized video is generated via an autoencoder from the source video, optical flow, and confidence map. To keep the utility of the sign language high, they use a loss function that especially focuses on the difference between hand and face motion of the driving and anonymized video. Again, there was no evaluation of whether the people could be identified by their hand motions.

5.3.4 Evaluation. Hand motion anonymization is mostly evaluated in the context of authentication and as such the **false positive rate (FPR)**, **false negative rate (FNR)**, and **equal error rate (EER)** are important metrics for evaluating the performance. But there is also the usage of

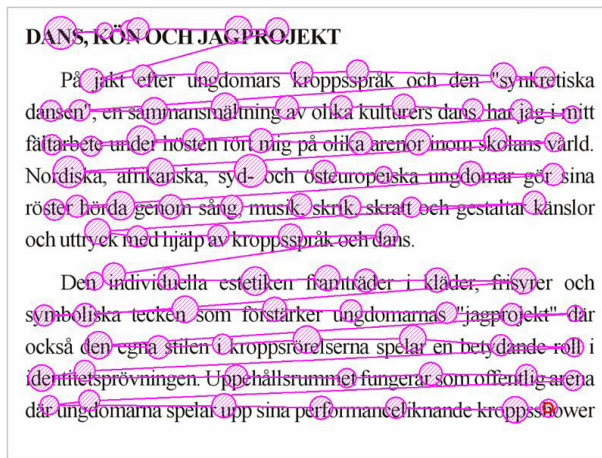


Fig. 7. Fixation and saccades while reading, from a study of speed-reading made by Humanistlaboratoriet, Lund University, in 2005. Source: <http://en.wikipedia.org/wiki/File:Rea>.

recognition approaches for the evaluation that uses the accuracy of identity, age, gender, and handedness inference. A unique evaluation approach we found was used by Goubaru et al. [97], who used the randomness of the template bits via occurrences and autocorrelation to evaluate their approach. Again, we find that more critical evaluation approaches are possible, as the EER will most likely overestimate the anonymization performance as it tries to achieve a low false positive rate.

5.4 Eye Gaze

Eye gaze involves two types of movements: **fixations** and **saccades**. Our eyes alternate between them during visual tasks, such as reading (see Figure 7). Fixations refer to maintained visual focus on a single stimulus, while saccades are rapid eye movements between fixations to reorient our gaze. Besides, even during fixations, our eyes are not completely still, but constantly producing involuntary micro movements (hundreds per second) known as microsaccades [4].

Eye-tracking technologies are becoming increasingly available in the consumer and research market. The most common type of tracking technology works by illuminating the eye with an array of non-visible light sources that generate a corneal reflection. These reflections are sensed and analyzed to extract eye rotation from changes in reflections. There is a wide range of hardware configurations for eye-tracking, including embedded cameras in computers, smartphones and virtual reality headsets, dedicated external hardware, or mobile eye-wear. These sensors allow to extract measurements not only regarding movement data related to fixations and saccades (speed, gaze angle, attention spots, scan path) but also additional features, such as pupil size variations and blink behavior. Combinations of these features provide valuable information to implement eye-gaze-driven applications.

5.4.1 Utility. Eye movements have been studied, analyzed, and used for more than a century in different research domains. In the medical field, gaze provides useful information about our cognitive and visual processing [16, 106], which can be used for diagnosing different diseases. In computer science, eye gaze is used as a form of human-computer interaction to improve accessibility, user experience, and to adapt system behavior [50, 169, 228]. More recently, security and privacy researchers have focused on analyzing stable unique features of eye movement to build biometric authentication systems [136]. Behavioral eye biometrics have been subject of intense investigation in the past decade, showing EERs as low as 1.8% [76]. Across all these

different domains, the utility to be preserved would depend on the underlying application, e.g., accuracy in predicting the next eye movement, in diagnosing a mental disease, in detecting the focus of user attention, or in recognizing a user.

5.4.2 Threat Space. Eye movement data is rich in information that can be exploited by malicious entities or curious service providers to uncover user-sensitive attributes beyond those disclosed intentionally and required for the purpose of the service or to directly identify a person. Besides the biometric information carried by eye movement data, research has also documented their correlation with multiple disorders and mental conditions, such as Alzheimer's [123], schizophrenia [116, 152], Parkinson's [148], bipolar disorder [89], mild cognitive impairment [304], multiple sclerosis [67], autism [31, 291], or psychosis [80], to name a few. Furthermore, pupil size is known to be an indicator of a person's interest in a scene [109] and a proxy for detecting cognitive load [146, 176]. Other recent works demonstrated that eye data can be used to infer gender and age or even personality traits [24, 147]. Given the richness of eye data and the increased availability of consumer tracking devices and the advent of eye-gaze-driven applications, there is a significant and imminent privacy threat potential [6]. The privacy threats of eye-tracking technologies have also been recognized by hardware makers like Apple, which disallow the usage of eye-tracking information for third-party applications in their Vision Pro Headset.

The two main threats that endanger eye privacy are re-identification and attributes' inference.

5.4.3 Anonymization Techniques. We found multiple recent proposals to protect the privacy of eye movement data, with many of them using noise injection to achieve **differential privacy (DP)**.

Random Perturbation. David-John et al. [55] adapt the task-based marginal model for eye gaze, in which for each feature vector dimension a distribution of the values is built to then randomly sample new synthetic data from these distributions. The identification accuracy of the generated synthetic data is close to chance level.

Noise Injection. Steil et al. [265] propose a DP-based technique to protect eye movement data collected while users read different types of documents (comic, newspaper, textbook) in a VR setting. The utility goal is to accurately predict the type of document to provide enhanced features in the reader application. Additionally, the privacy goals are to avoid gender inferences from eye movement data and to protect against re-identification when the attacker has prior knowledge of a dataset including the target user's eye data and identity. To achieve these goals, the exponential mechanism [74] is applied to a database of users' eye features by a trusted curator prior to its release. This sanitized database can be then used for training classifiers to provide the enhanced reader functionality. The experiments testing at various noise levels shows that utility with regard to document classification can be partly preserved (~55%–70%) while reducing gender accuracy inference to the level of random guesses (~50%).

Based on Steil et al.'s dataset, Bozkir et al. [33] evaluate two types of DP-based perturbations: the standard **Laplacian perturbation algorithm (LPA)** [73] and the **Fourier perturbation algorithm (FPA)** [239]. They also propose a modification of the FPA algorithm that splits eye data in chunks before adding noise to reduce temporal correlations, which is a source of reduced utility, as more noise is required to protect privacy. With this modification, they obtain document type classification results similar to those used by Steil et al. [265] for the case of 50% gender classification, while adding more noise to the data (better privacy guarantee).

Liu et al. [156] present a DP-based solution to anonymize eye tracking data aggregated as a heatmap. A heatmap, or attentional landscape, is a popular method for visualizing eye movement data that represents aggregate fixations [72]. This means that the intensity of every pixel is

adjusted relative to the number of fixations over that region. The privacy goal in this case is to protect individual gaze maps while preserving the utility of the aggregated heatmap. Their experiments with random selection and additive noise (Gaussian, Laplacian) show that Gaussian noise is the best option to obtain good privacy guarantees for the individuals' gaze maps without visually distorting the hotspots in the aggregated heatmap, i.e., keeping a certain utility.

David-John et al. [57] worked on protecting eye tracking data recorded in VR/AR headsets. They propose two different interface models for how data can be shared with a third party and propose three anonymization techniques—Gaussian noise injection, temporal down-sampling, and spatial down-sampling—for one of the interface models. The noise injection approach was found to be the most effective, as it reduced the identification rate of the subjects the most with high variance values for the Gaussian distribution. Wilson et al. [295] also proposed adding Gaussian noise to eye tracking data, showing similar results.

Hu et al. [119] proposed a local differential private mechanism for generating synthetic eye movement trajectories called Otus. Their technique first separates the field of view into tiles and then constructs a graph that encodes the gaze duration of each tile and the transition probability between the tiles. The graph is then perturbed using the Laplacian mechanism before it is sent to the server. The server then averages all user graphs and uses random walks on the graph to generate new eye movement trajectories.

Li et al. [153] proposed Kaleido, a plugin system that can be used to anonymize eye gaze trajectories with differential privacy guarantees. The authors extend geo-indistinguishability [14] and w-event privacy [137] to take into account the area of interest with radius r a user is looking at. The intuition of their guarantee is that all gaze positions within the area are indistinguishable. They note that they only protect against spatial information and not temporal information. Further, they define an adaptive algorithm to allocate the privacy budget of a user, depending on the total privacy budget of each time window. Their results show a reduction of the identification of users to near chance level, however, the utility of the data is also close to chance level.

Coarsening. The temporal and spatial down-sampling proposed techniques by David-John et al. [57] are both coarsening-based techniques. For the temporal down-sampling, only a very small reduction in the identification accuracy can be recorded while the spatial down-sampling has a bigger effect but must be scaled very high to do so. Wilson et al. [295] proposed a spatial down-sampling approach for the eye gaze angle. They first map the 180° to 2,160 points and then coarsen the gaze angle to these points. In their evaluation, the spatial down-sampling seems to be more effective than temporal down-sampling.

Continuous Conversion. Wilson et al. [295] propose smoothing the eye gaze using a sliding window approach. They show that using a large enough window reduces the identification rate.

David-John et al. [55] applied k -anonymity to eye movements by grouping the trajectories of users and then averaging them. They were able to show that even with small numbers of k the identification accuracy drops significantly. Due to them processing the feature vectors of each task separately, their reported high utility is questionable. In a follow-up paper, David-John et al. [56] propose two synthetic data generation approaches for eye gaze. Their k -same synth approach applies k -anonymity to the fitted parameters of a Gaussian mixture model before using it to generate fixations and saccades. Their event-synth-PD approach uses a conditional variational autoencoder to generate new data with given characteristics. They show that their event-synth-PD approach achieves plausible deniability. They compare both methods to Kaleido and achieve comparable results for privacy and utility.

Fuhl et al. [86] perform eye gaze anonymization by using an autoencoder in combination with reinforcement learning. The autoencoder is trained on the eye gaze trajectories to learn a latent

representation of the data. Then, a manipulation agent modifies the latent vector of the trajectories to prevent, for example, gender classification. After the decoding of the latent vector, a classifier tests how good the manipulation was, and its result is used as the loss for the training of the manipulation agent.

5.4.4 Evaluation. The proposals by Steil et al. [265] and Bozkir et al. [33], measure the quality of their anonymization techniques for attribute inference protection using the classification accuracy metric for the main task and the attribute inference task. For the re-identification protection case, it is assumed that the attacker has previous knowledge of a database of users' eye data and their identities. To simulate this knowledge, they train the classifiers on the clean data and test them on the anonymized data using also the accuracy metric to report privacy protection. Besides, these works also report the so-called privacy loss parameter (or ϵ) from DP theory, which quantifies the maximum difference between the data points of two individuals in the dataset.

Liu et al. [156] analyzed the privacy-utility tradeoff of anonymized heatmaps using the **correlation coefficient (CC)** and **mean square error (MSE)** of noisy heatmaps under different privacy levels (different values of ϵ). The CC and MSE give an idea of the similarity between the original and the anonymized heatmaps, and the ϵ provides information about the privacy guarantee (the smaller, the better privacy). These metrics are accompanied by the visual representation of the noisy heatmap to aid the relevant stakeholders in deciding what level of noise is acceptable for a given application.

Regarding datasets, the largest dataset available is GazeBaseVR [160], which captured 407 participants performing five tasks with up to six sessions. As recording device, they used a VR headset. Steil et al. [265] collect data from 20 participants (10 male, 10 female, aged 21–45) while reading documents using a VR headset. Each recording is divided into three sessions (reading a comic, newspaper, or textbook) lasting 30 minutes in total. They extract 52 eye movement features related to fixations, saccades, blinks, and pupil diameter. The dataset has been publicly released⁴ by the authors, and Bozkir et al. [33] use it as the basis to evaluate their proposal.

The Ehtask [120] dataset contains the recordings of 30 people performing four different eye gaze tasks using a VR headset. Another VR headset dataset is DGaze [71], which captures 43 people in five different scenes. In the heatmaps anonymization study, Liu et al. use a synthetic simulated dataset to illustrate their privacy analysis. Besides the technical privacy analysis, Steil et al. [265] is one of the few works considering user privacy concerns regarding behavioral data collection. They conduct a large scale user survey (with N=164 participants) to explore with whom, for which services, and to what extent users are willing to share their gaze data. Their report shows that people are uncomfortable with inferences (gender, race, sexual orientation) and would object to sharing their data if these attributes can be leaked. The results also show that people generally agree to share their eye tracking data with a governmental health agency or for research purposes but would object to doing so if the data owners are companies. These insights are a first step towards understanding user privacy awareness and privacy needs, but more work is required in this field to guide the design of user-centered privacy protective techniques for behavioral data.

5.5 Heartbeat

An **electrocardiogram (ECG)** is a graph of voltage over time that captures the electrical activities of cardiac muscle depolarization followed by repolarization during each heartbeat. Shown in Figure 8, the ECG graph of a normal beat is composed of a sequence of waves: a P-wave reflecting

⁴<https://www.mpi-inf.mpg.de/departments/computer-vision-and-machine-learning/research/visual-privacy/privacy-aware-eye-tracking-using-differential-privacy>

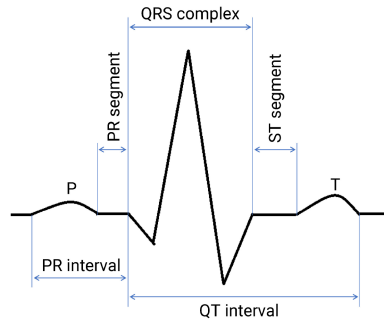


Fig. 8. Waveform of an ECG signal with normal cardiac cycle. Source: https://www.nottingham.ac.uk/nursing/practice/resources/cardiology/function/normal_duration.php.

the atrial depolarization process, a QRS complex representing the ventricular depolarization process, and a T-wave denoting the ventricular repolarization. Other portions of the ECG signal encompass the PR, ST, and QT intervals [323].

Like other biometric systems applied to identification tasks, ECGs are typically converted into abstract, compressed representations, typically referred to as biometric templates, before the task is conducted. Biometric-template methods can be classified depending on the exploited features of the ECG data. The most popular ones are fiducial-based, non-fiducial-based, and hybrid methods [207]. On the one hand, fiducial-based techniques utilize characteristic points on the ECG signal to extract temporal, amplitude, envelope, slope, and area features. Characteristic points are the locations that correspond to the peaks and boundaries of the P, QRS, and T-waves of the ECG signal. On the other hand, the non-fiducial-based methods do not rely on the ECG characteristic points, and examples include autocorrelation coefficients, Fourier and wavelet transforms. Hybrid methods combine both fiducial-based and non-fiducial-based features.

5.5.1 Utility. ECG data find application in healthcare and biometrics systems, the latter being intended for identification and authentication [284]. In healthcare, ECGs are utilized for diagnosis of heart diseases [158]. Typically, there is a stand-alone service or a complete e-health system where the service provider, in addition to offering a repository of personal medical data, may allow to remotely process such data. In any case, the aim is to provide real-time feedback to patients and hospitals, either as a warning of impending medical emergency or as a monitoring aid during physical exercises.

5.5.2 Threat Space. Regardless of the application (i.e., identification, authentication, or healthcare), ECGs are health data and, as such, are considered sensitive by data-protection regulations and need to be protected. Consider the case, for example, of a user who might see their insurance premium increased or suffer discrimination during a job application due to a medical condition inferred from their ECGs.

Although it is well known that ECG data may help diagnose a patient's physiological or pathological condition, other probably lesser-known inferences include cocaine use [118] and stress [224], which may be sensitive to the patient and obviously should be kept private. The fact that the very same time series data allows drawing both desirable inferences (i.e., for healthcare) and sensitive inferences (that need to be protected) poses a dilemma of great practical relevance.

5.5.3 Anonymization Techniques. Next, we survey the most relevant privacy-protection techniques for ECG data.

Another approach based on **compressive sensing (CS)** [38] is proposed by Djelouat et al in Reference [70]. CS is a signal processing technique that combines both sampling and compression through random projections. Building on this technique, the authors propose compressing the ECG signal by sampling it at the time of sensing. This reduces the need to even store the sensitive ECG data at the wearable device, thereby providing protection against that entity. The theoretical properties of this compression technique ensure that, under certain assumptions on the random projection, a good reconstruction of the original ECG signal can be obtained at the provider side.

Feature Removal. Kalai et al. [317] present a template protection scheme for ECG data. In a first phase, the authors propose computing the **discrete cosine transform (DCT)** of the ECG signal's autocorrelation coefficients and then removing those DCT coefficients with the lowest energy. The remaining DCT coefficients constitute the biometric template. In a second phase, two keys are obtained from the template. One is transmitted to the target application the user wishes to authenticate. The other functions as a private key, which is derived from the complete DCT already stored in the server. A similar approach is presented by Zaghouni et al. [318] that uses a quantization step once the DCT-template is obtained. This latter approach is evaluated on the PTB dataset, but no experimental comparison is conducted between the two proposed solutions.

Another similar proposal is made by Mahmoud et al. [166] that decomposes the ECG signal into its wavelet transform, eliminates the low-frequency coefficients, and reconstructs the ECG signal for release. At the provider side, only authorized personnel with access to a secret key (derived from the wavelet-transform template) is able to reconstruct the original ECG from the released, protected signal. To what extent these released data may safeguard patients' privacy is evaluated through the **percentage root mean square difference (PRD)**, a simple and widely used distortion measure in ECG signal processing applications [172] that quantifies the difference between the original ECG and its protected version.

Continuous Conversion. Bennis et al. [23] proposed a simple k-anonymity scheme for ECG data. In their first step, they transform the signal into the frequency domain. Next, they pick the k closest neighbors of the signal and then aggregate those into a new signal before transforming it back into the time domain.

Piacentino et al. [222] used a GAN to generate synthetic ECG data by first normalizing the data and then arranging it into a matrix. For the arranging of the data, multiple proposals are made sorting the data values by their type. No evaluation of the privacy of the synthetic data was performed. Jafarlou et al. [126] also propose to use a GAN to generate anonymized ECG data samples. Their approach differs from Piacentino et al. in that they use the original ECG sequence as input to the GAN and use the identification accuracy as part of the training loss for the GAN. Their evaluation shows lower identification accuracies while still allowing arrhythmia detection. Nolin-Lapalme et al. [204] also use a GAN for the ECG anonymization, but they aim at generating sex-neutral ECG samples and use the sex classification as part of the GAN loss.

Random Perturbation + Noise Injection. Although encryption based on the idea of CS can achieve a computational notion of secrecy through the random projection step, it has been shown this technique is vulnerable from an information-theoretic perspective [238]. To address this problem, Chou et al. [46] propose using principal component analysis and SVD on a CS scheme, where the ECG data is encrypted at the wearable sensor by adding signal-dependent noise. They measure privacy as the mutual information between the original ECG signal and its encrypted version and show that high classification accuracy can be achieved while providing privacy beyond computational secrecy.

Discrete Conversion + Noise Injection. Unlike the works surveyed previously, the goal of Zare-Mirakabad et al. [319] is to publish suitable representations of ECG data with certain privacy guarantees. To do this, Zare-Mirakabad et al. propose converting ECG time series into symbolic representations over time. They use the popular **Symbolic Aggregate approXimation (SAX)** to replace continuous numerical values with strings of symbols. With this new symbol representation, the proposed anonymization technique first builds an n -gram model from the complete time-series string and then ensures that each n -gram has a minimum frequency of occurrence, similar to the k -anonymity criterion. To ensure this version of k -anonymity is satisfied over the string of symbols, the authors contemplate adding fake n -grams to the original string. Experimental results on the Eamonn Discord Dataset show that (a measure of) information loss is hardly affected for values of k up to 20.

Continuous Conversion + Random Perturbation. Chen et al. [44] and subsequent work by Wu et al. [296] address the problem of making ECG-based biometric templates revocable, just like keys or passwords, a property they consider indispensable for ECGs to be used in practice. To enable template revocability, the common practice is to associate distinct templates with the same biometrics by perturbing them in a different manner. To protect user privacy, however, this process needs to ensure the recovery of the original biometric from its template is either infeasible or computationally hard.

Essentially, cancelable templates are obtained as random projections of a user's ECG data block. Unlike common approaches, however, Wu et al. put no restrictions on the generator matrix. Accordingly, the idea is that each realization of this matrix allows canceling their corresponding templates. Reidentification is then conducted with the multiple-signal classification algorithm [26], reporting rates of over 95% in the Physikalisch Technische Bundesanstalt Database.

A distinct approach by Hong et al. [117] proposes a template-free identification system to prevent any privacy issue from compromised or stolen templates. The system converts ECG-data into images through various spatial and temporal correlations methods and uses deep-learning techniques to train a classifier. The authors conduct experiments on the Physikalisch-Technische Bundesanstalt database and report identification rates of over 90% with sampling rates of 1,000 Hz.

Continuous Conversion + Noise Injection. Sufi et al. [268] propose building templates of the waves P, QRS, and T through cross-correlations of the ECG signal. Each of those templates are then obfuscated in a concatenated fashion with additive noise generated synthetically, so the obfuscation of a wave serves as input to obfuscate the next wave. The results are noisy forms of the three waves and noisy templates thereof. All this information constitutes the key available to authorized personnel, who will be able to reconstruct the original ECG from the noisy version (which is shared or made publicly available by the patient or user themselves). Unauthorized personnel, will only have access to the noisy ECG signal, which, according to the authors, may prevent identity and attribute disclosure.

Huang et al. [121] propose an authentication system that protects the privacy of ECG templates in a database with differential privacy. The authors assume the interactive setting of this privacy notion, where an analyst queries the database to obtain ECG data. Specifically, the analyst is supposed to ask for the coefficients of a Legendre polynomial, which the anonymization system utilizes to fit and compress the ECG signal. Laplace noise is calibrated to the sensitivity of those coefficients and added to them, and the noisy response is returned to the analyst. The ϵ parameter of DP therefore regulates the tradeoff between user privacy and authentication accuracy, the latter aspect depending on two sources of error: the polynomial fitting approximation and the injected noise. The authors evaluate the system in the MIT-BIH ECG and MIT-BIH Noise Stress databases, reporting decent authentication accuracy. However, they appear to misunderstand how the

sensitivity of the coefficients is computed, and therefore their results seem to have been obtained incorrectly.

Saleheen et al. [249] investigate if sensitive inferences from segments of time series data can be drawn by a dynamic Bayesian network adversary. The adversary is assumed to estimate a range of behavioral states about the user, including, for example, whether or not they are in a conversation, running, smoking or stressing, at the time the data is gathered. When the adversary is likely to infer sensitive aspects of a user, the corresponding segments of data are substituted for most-plausible, non-sensitive data. To estimate the privacy provided by these substitutions of data, the authors propose a variation of the differential-privacy notion that bounds the information leaked resulting from the substitutions. In other words, the proposed metric ensures that the information leaked about a sensitive inference from a substituted segment is always bounded. Utility loss is, however, computed as the absolute difference between the probability of inference about each non-sensitive behavioral state from actual data and the same probability from released data. Although experimental results show relatively small values of utility loss for $\epsilon \in [0.05, 0.65]$, the proposed solution has two main limitations: First, protection is provided only for dynamic Bayesian network adversaries; and, second, it assumes all time-series data are available beforehand, which precludes its application in real-time scenarios.

5.5.4 Evaluation. The reviewed techniques measure how service functionality is degraded due to anonymization with common machine learning metrics such as precision, recall, and accuracy, and less frequently with the DTW and PRD quantities, which assess the similarity between original and protected time series. As for privacy, the level of protection is assessed through a variety of notions and measures, including the accuracy of a membership inference attack, the ϵ parameter of differential privacy, the mutual information between the original ECG signal and its encrypted version, the probability of correct inferences on sensitive attributes with and without protection, and through a notion similar to k -anonymity. A common dataset used is the MIT-BIH arrhythmia database [190], which contains the ECG samples of 47 people.

5.6 Brain Activity

Brainwaves are patterns of measurable electrical impulses emitted as a result of the interaction of billions of neurons inside the human brain. Since the first human electroencephalogram was recorded in 1924 [101], both the hardware devices to measure brain activity and the analysis techniques to process these signals have significantly improved. Current technologies to measure brainwaves can be classified as invasive and noninvasive methods. Invasive methods record signals within the cortex by directly implanting electrodes near the surface of the brain [133]. These methods are far too risky for usage under noncritical circumstances and are only used in clinical applications. Instead, non-invasive methods are most frequently used and applicable to many areas other than the medical realm, such as brain-controlled interfaces. The most portable and commonly used of these techniques is EEG, which records electrical activity through sensors placed on the scalp surface.

An EEG signal is a combination of different brainwaves occurring at different frequencies. Every type of wave carries different kinds of information, which can be used to gain insights about the current state of the brain [10]. Researchers have tried to identify certain mental states associated to each brainwave. Table 1 presents a summary of the most important wave types, their respective frequencies, their originating location in the brain, and their associated mental state.

Brain-computer interface (BCI) technologies mostly work on continuous EEG data recordings, i.e., time series data. But there are also many applications based on the extraction of time-locked brain variations that appear in reaction to external stimuli. These variations, called

Table 1. Overview of EEG Brainwaves—Based on References [10] and [2]

Wave Type	Freq. (Hz)	Originating Location	Mental State
<i>Gamma</i> γ	30-100	Somatosensory cortex	Active information processing, strong response to visual stimuli [2]
<i>Beta</i> β	13-30	Both hemispheres, frontal lobe	Increased alertness, anxious thinking, focused attention
<i>Alpha</i> α	8-13	Posterior regions, both hemispheres; High amplitude waves	Resting, eyes closed, no attention [139]; Most dominant rhythm
<i>Theta</i> θ	4-8	No special location	Idling, dreaming, imagining, quiet focus, memory retrieval
<i>Delta</i> δ	0.5-4	Frontal regions; High amplitude waves	Dreamless and deep sleep, unconsciousness

event-related potentials (ERPs), are widely used to detect neurological diseases. In both cases, either using ERPs or a longer EEG series, features are computed for the brainwave data-driven application built on top. These features can belong to the time and/or frequency domain and to one or multiple channels. Examples of commonly used features include Autoregressive coefficients, Fourier and Wavelet transforms.

5.6.1 Utility. The utility that should be preserved when processing brainwave data is highly dependent on the application. For clinical applications, for example, the raw information could be needed for a proper diagnosis or a safe brain-controlled prosthesis. In these cases, regulations like the HIPAA Privacy Rule [111] are usually in place to protect personal identifiable information. When moving to other less-regulated fields of application, the need for full raw EEG data is not necessarily justified. The most prominent EEG applications include user authentication, personalization of gaming experiences, and brain-controlled interfaces. In these cases, the utility to be preserved should be enough to provide a useful application, i.e., recognize the user, and offer personalized options and responsive interfaces all with a tolerable error that does not hamper the security and usability of the service.

5.6.2 Threat Space. Brain activity is rich in information. It can be used to uniquely identify individuals given their unique characteristics and, in fact, several biometric systems based on brainwaves have been proposed [99]. Besides, the acquisition of EEG signals raises privacy issues, because brainwaves correlate, among others, with our mental states, cognitive abilities, and medical conditions [269]. Martinovic et al. [173] demonstrated that by manipulating the images presented to the users, their EEG signals could reveal private information, e.g., bank cards, PIN numbers, area of living, or if the user knew a particular person.

5.6.3 Anonymization Techniques. We found that a large number of anonymizations rely on machine learning methods to perform the anonymization of the data, with approaches like **Generative Adversarial Networks (GANs)** and adversarial perturbation scheme dominating the field. With the availability of EEG datasets, the anonymization of brain activity data is gaining some traction.

Feature removal. Matovu et al. [174] explore how to reduce the leakage of private information from EEG user authentication templates. They assume an insider type of attacker, such as an unscrupulous database administrator, who misuses their privilege to maliciously exploit the templates. The attacker wants to infer, specifically, if the user associated with a template is an alcoholic. Their envisioned anonymization technique aims at concealing the alcoholism information

while still providing good authentication accuracy. It is, therefore, an attribute protection mechanism. Conceptually, it is based on the hypothesis that different template designs (features, channels, frequencies) will have an impact on the amount of non-authentication information (emotions, health conditions) that can be inferred. The authors demonstrate this hypothesis by choosing two different templates and calculating the predictive capability to authenticate users and determine their alcohol consumption behavior.

Continuous Conversion. In the same direction of feature selection, Yao et al. [310] propose the usage of GANs [96] to filter sensitive information out of EEG data. Their goal is to reduce the possibility of inferring alcoholism while keeping the brain activity recordings useful to detect mental tasks—specifically, to predict which visual stimulus the user is looking at. The GAN-based proposed filter involves deep neural networks that perform domain transformation, that is, translating EEGs from a source domain distribution X with both desired and privacy-related features to a target domain distribution Y with desired features only. Their results after applying the filtering technique show a significant reduction in the percentage of EEG sequences from alcoholic users that can be classified as such (from 90.6% to 0.6%). At the same time, the mental task classification accuracy does not drop significantly (4.2% less). However, the original mental task classifier accuracy was not strong before filtering the privacy-sensitive features, and it remains to be studied if this technique would work in other classification scenarios.

Pascual et al. [216] use a GAN to generate synthetic EEG data to train an epilepsy monitoring system, as sharing large amounts of medical EEG is a privacy problem. The authors focus on inter-ictal EEG signals (signals between two seizures), as these are easier to record than the actual seizures. As generator, a convolutional autoencoder is used, but instead of decoding an inter-ictal, the latent code is translated into an ictal sample. The discriminator then compares the synthetic ictal to a real one. Their results show that the synthetic data reaches identification rates that are close to chance level, even when only two patients are in the test set. However, this is only a pseudonymization of the patients, as all synthetic ictal values generated for a specific patient can still be linked to each other.

Bethge et al. [25] proposed privacy encoders to remove the sensitive information from each of the brain activity data streams before they are used in a classification task. For each dataset, a convolutional neural network is trained as encoder using the **maximum mean discrepancy (MMD)** between the different encoded datasets as loss function. This way the encoders should learn a domain-invariant representation of the data. They test their approach on four datasets, finding that the classification from which dataset a sample originated drops from 99% to 52%, while the emotion classification is only reduced from 51% to 49%. It remains an open question how well the identity of a subject would be preserved by this approach. A similar approach is being proposed by Meng et al. [181]; instead of using a neural network for the transformation, they learn a perturbation vector that is added to the EEG signal. The perturbation is learned via an adversarial scheme using an action classifier to establish the utility and a biometric recognition system for the privacy. Another adversarial approach is being proposed by Singh et al. [258]. The main difference from the previous approaches is that an autoencoder is used for the transformation.

Continuous Conversion + Noise injection. Debie et al. [62] also use a GAN to generate new synthetic data from the original one. They differ from Yao et al. and Pascual et al. in that they use differentially private stochastic gradient descent on the discriminator of the network. This method reduces the influence of each individual to the computation of the gradients. They evaluated their GAN on the Graz dataset A with EEG data from nine subjects. Their results show that the utility of the synthetic data is well preserved, however, no additional privacy evaluation was performed.

5.6.4 Evaluation. The reviewed works, similar to the proposals for anonymizing gait, evaluate the quality of inference protection by comparing the prediction accuracy for the protected attribute before and after modifying the EEG data. The metrics used for this analysis are typical machine learning metrics, including accuracy, false positive rates, and false negative rates. Similarly, the loss of utility is evaluated by measuring the reduction in classification accuracy when using the original and anonymized EEG data.

For their evaluations, the works use a variety of different EEG datasets. The largest dataset is the Temple University Hospital EEG data corpus [206], which contains 579 subjects, followed by the BCI2000 dataset [251] with 106 subjects. Specifically recorded for authentication was the dataset of Arias et al. [15], which recorded 56 people. A special dataset is the SUNY medical dataset, with EEG data of 25 alcoholic subjects and 25 control subjects while looking at visual stimuli [134, 203]. Further, there exist a couple of smaller datasets [114, 266, 293].

6 Discussion

All reviewed behavioral biometric traits have in common that they are captured as a time-series tracking the change of the trait over time. Most traits, such as gait, hand motions, voice, and eye gaze are overt traits that can be observed from a distance and do not require the participation of the subject. These traits are often captured as a byproduct for other recordings, for example, video recordings. EEG and ECG, however, are secret traits that can mostly only be recorded by directly attaching sensors to the subject to measure them. We found the most anonymization methods for voice and the least for EEG. For the traits touch, thermal, and lip-facial, we could not find any mechanisms.

The **utility** of these traits is very diverse and is mostly unique to each trait and the application using it. It ranges from utilities such as the naturalness of a motion to the intelligibility of utterances.

Regarding their **threat space**, the traits are similar to each other, as due to the pervasiveness of digital capturing devices, more instances of them are captured. Wearables and mobile devices are of special interest, as they are attached to the subject and can therefore allow continuous capture of behavioral data. As our literature review has shown, all traits can be used for both identity and attribute inference, which then can be abused for a wide variety of privacy threats such as surveillance, identity theft, or private attribute inference. The privacy goals, identity protection, and attribute protection are also the same for all the traits. However, voice has an additional privacy goal in which the content of the speech should be made unintelligible.

For the **techniques** (see Table 2 and Table 3) that we reviewed, we found that most of them fall into the category of continuous conversion, followed by feature removal and noise injection. Next are random perturbation and discrete conversion, with most discrete conversion methods aiming at template protection. Coarsening is the category with the least amount of methods. We observe several differences for the categories of our taxonomy. For the removal methods, we find that the removal is not directly reversible, however, due to the high redundancy in behavioral biometric data, it still might be possible to reconstruct the removed data. For the conversion methods, we often observe that the parameter space for the anonymizations is often rather small, making it possible that an attacker can link clear and anonymized data by brute-forcing the parameters when the anonymization technique is known. In general, we find that the reversibility of conversion techniques still has to be evaluated better. For noise injection techniques, we find that the strong dependency both temporal and physiological features is a problem, since they can be used to filter out the noise.

With regard to the techniques providing **differential privacy**, we have observed that none of them can be used continuously over time without completely compromising user privacy.

Table 2. An Overview of All Found Methods Classified by Trait and Method

Trait Method	Voice	Gait	Hand motion	Eye-Gaze	Heartbeat	Brain activity
random perturbation	[193, 214, 254]	[113]	[97, 167, 168, 288]	[55]	[46]*	
noise injection	[103, 104, 107, 163, 208, 271, 285]	[105, 175, 183, 276, 277]	[188, 252]	[57, 119, 153, 156, 265, 295]		
coarsening		[198]	[168, 288]	[57, 295]		
feature removal	[201, 213, 215, 297, 322] [7, 49, 69, 200]	[63, 90, 105, 131, 246]			[166, 317, 318]	[174] [310]
discrete conversion	[27, 217, 229, 230]		[85, 151, 189, 194, 248, 288]		[319]*	
continuous conversion	[1, 3, 17, 82, 83, 129, 132, 138, 161, 164, 225, 234, 262] [9, 11, 45, 100, 110, 177, 186, 210, 218, 232, 263, 314] [306] [40, 184, 185, 202, 308] [162, 186, 187, 219, 307, 316] [43, 47, 65, 178, 240, 299, 309] [92, 233] [†] [257] [†] [253]*[39] [†] [143]* [144]*[236]*[235]*[263]*	[8, 102, 125, 275] [191] [112, 196] [‡]	[81, 167, 171, 250, 298]	[55, 86, 295] [56]	[23, 126, 222] [44, 204] [†] [296] [†] [117] [†] [268]* [121]*[249]*	[25, 181, 216] [62, 258]*

Papers that propose multiple methods can appear in multiple rows. Papers that combine multiple methods are marked the following: * plus noise injection, [†] plus random perturbation, [‡] plus discrete conversion.

Table 3. An Overview over which Privacy Goals the Different Techniques Try to Achieve

Trait Privacy Goal	Voice	Gait	Hand motion	Eye- Gaze	Heartbeat	Brain activity
Attribute	[11, 27, 39, 47, 83, 103, 110, 163, 193, 217, 229, 230, 234]	[90, 105, 113]	[97, 167] [168, 288] [189, 248] [167]	[33, 86, 265]	[44, 117, 121, 166, 249, 268, 296, 317, 318]	[25, 62, 174, 310]
Identity	[1, 3, 7, 9, 11, 17, 45, 49, 82, 100, 103, 104, 107, 129, 132, 138, 143, 144, 161, 163, 164, 177, 200, 201, 208, 210, 213–215, 218, 225, 232, 234–236, 254, 262, 263, 285, 297, 314, 322] [110, 185, 186, 202, 306, 308] [40, 162, 184, 187, 307, 316] [43, 65, 186, 219, 309] [92, 178, 233, 240, 257, 299] [253]	[8, 63, 90, 102, 105, 112, 125, 131, 175, 183, 198, 246, 275–278] [191, 196]	[81, 85, 151, 171, 188, 194, 250, 252, 298]	[33, 55, 55, 57, 86, 119, 153, 156, 265, 295] [295, 295] [56]	[23, 46, 126, 222, 319] [204]	[62, 181, 216] [258]

The reason lies in that the privacy budget is necessarily finite, which means, by the sequential composition property of differential privacy [179], that it will be consumed completely at some time instant. Surprisingly, this appears to be in contradiction to the intended use of most of the applications where differential privacy is guaranteed, namely, continuous monitoring in health-care scenarios and identification and authentication services (which clearly are not single-use services). In that respect, the use of related privacy notions intended for continuous observations (e.g., w -event differential privacy [137]) may come in handy. In general, more research is needed on how to effectively apply differential privacy to behavioral data.

We made the observation that most methods do not **manipulate the temporal aspect** of their data. Notable exceptions are Hirose et al. [112] and Maiti et al. [168]. Since all traits result in time series data, manipulating the temporal order or time differences between events could lead to some general anonymization techniques that work for multiple traits. For attribute protection, we find anonymizing intrinsic attributes (e.g., age, sex) to be difficult, as it is not clear which part of the behavioral data is relevant for these attributes. We therefore find generative machine learning approaches a promising approach to address this problem, as the machine learning models can learn the intrinsic dependencies between data and attributes. Further, we noticed a lack of even a basic understanding of **users' privacy awareness** and concerns about behavioral privacy. These are necessary to design protection techniques that consider user needs and requirements.

We found that the **evaluation methodology** between the traits and methods is rather similar. In general, an inference/recognition system is being used on the clear and on the anonymized data and then the difference in accuracy is reported, often without retraining the inference system on the anonymized data. We find this methodology too simple, as the underlying assumption is that the attacker is not aware of the anonymization. A notable exception are more recent voice anonymization techniques that now mostly rely on the benchmarking framework of the VoicePrivacy Challenge to evaluate the privacy and utility of their techniques. This shows that community initiatives can provide a common basis for comparison and improve the overall evaluation methodology of a field.

Only a small number of articles compare their own methods to that of others, and due to the differences in attacker models and data sources, they are difficult to compare for the readers. We also found that there are not many approaches [237, 321] to formalize the privacy of behavioral biometric anonymization methods, and most of the evaluations rely on empirical privacy estimations. Another problem is that the evaluation methodology is too close to the recognition system evaluation methodology that seeks to infer persons in a large dataset with poor data quality, while an anonymization method should also work on a small group size with high data quality. We believe that the lack of available datasets (see Table 4) is one of the main problems that keeps the less-researched behavioral biometric traits back. For possible future work, we see the anonymization of eye-gaze and motion data as promising areas of research, as many challenges remain, such as achieving good utility and real-time applicability. Similar to the VoicePrivacy Challenge, most behavioral biometrics would benefit from community-driven evaluation frameworks to increase the comparability and rigor of privacy and utility evaluations. One area where many behavioral biometric traits are combined is the creation of digital twins, where it is an open question whether anonymizing the behavioral traits independently of each other is sufficient to create privacy-friendly digital twins, e.g., for mixed reality.

7 Concluding Remarks

Anonymizing behavioral biometric data is an important task for protecting people's privacy. In our literature review, we found many different behavioral traits that need to be considered and developed a taxonomy to classify the anonymization techniques that can be applied to them by

Table 4. An Overview of Used Behavioral Biometric Datasets

Name	Participants	Published	Source	Trait
TIMIT	630	1993	[91]	Voice
Albayzin	164	1993	[192]	Voice
YOHO	137	1994	[37]	Voice
BioSecureID	400	2009	[84]	Voice
Billeb et al.	701	2014	[27]	Voice
Librispeech	1,166	2015	[211]	Voice
RSR2015	300	2015	[150]	Voice
VCC 2016	10	2016	[280]	Voice
DAIC-WOZ	189	2016	[286]	Voice
VoxCeleb	1,251	2018	[195]	Voice
CSTR VCTK Corpus	110	2019	[302]	Voice
AISHELL-3	218	2020	[311]	Voice
Kassel State of Fluency	37	2022	[20]	Voice
CASIA-B	124	2005	[325]	Gait
BEHAVE	125	2010	[28]	Gait
OU-ISIR	200	2012	[170]	Gait
EPIC-Kitchens	32	2020	[52]	Gait
IITMD-WFP	31	2021	[274]	Gait
ETRI-activity 3D	100	2020	[127]	Motion
NTU60	40	2020	[157]	Motion
BOXRR-23	105,852	2023	[197]	Motion
MCYT baseline corpus	330	2003	[209]	Hand motion
SVC2004	100	2004	[313]	Hand motion
GREYC	133	2009	[93]	Hand motion
MNIST	500	2012	[66]	Hand motion
Web-based keystroke	83	2012	[94]	Hand motion
SMILE	30	2018	[77]	Hand motion
ASLLRP	33	2022	[199]	Hand motion
DOVES	29	2009	[287]	Eye-Gaze
VR-Saliency	169	2018	[259]	Eye-Gaze
Gaze Prediction	43	2018	[300]	Eye-Gaze
Video viewing	50	2017	[159]	Eye-Gaze
MPIIDPEye	20	2019	[265]	Eye-Gaze
OpenEDS	157	2019	[88]	Eye-Gaze
EHTask	30	2022	[120]	Eye-Gaze
DGaze	22	2020	[71]	Eye-Gaze
GazeBaseVR	407	2023	[160]	Eye-Gaze
SUNY EEG database	50	1999	[203]	Brain activity
UCI EEG database	122	1999	[21]	Brain activity
BCI2000	106	2004	[251]	Brain activity
DEAP	32	2011	[142]	Brain activity
SEED	15	2015	[326]	Brain activity
DREAMER	23	2018	[135]	Brain activity
Temple University Hospital	579	2016	[206]	Brain activity
Arias et al.	56	2021	[15]	Brain activity
MIT-BIH ECG Arrhythmia	47	1979	[190]	Heartbeat
Phys. Technische Bundesanstalt	290	1995	[32]	Heartbeat

the type of data transformation they perform. While voice anonymization is already an established research field with many insights, most behavioral biometric traits only received little attention. Their protection hence remains an open research question. We further found that most anonymization techniques are only evaluated rudimentarily with the assumption of a weak attacker. Improving the evaluation methodology is therefore another open research question. Last, we find that the temporal aspect of the data was mostly neglected: On the one hand, only few anonymization approaches exist for data streams, and on the other hand, most anonymization techniques do not perturb the temporal aspect of their data.

References

- [1] Alberto Abad, Alfonso Ortega, António Teixeira, Carmen García Mateo, Carlos D. Martínez Hinarejos, Fernando Perdigão, Fernando Batista, and Nuno Mamede (Eds.). 2016. *Advances in Speech and Language Technologies for Iberian Languages (Lecture Notes in Computer Science, Vol. 10077)*. Springer International Publishing. DOI : <https://doi.org/10.1007/978-3-319-49169-1>
- [2] Mohammed Abo-Zahhad, Sabah Mohammed Ahmed, and Sherif Nagib Abbas. 2015. State-of-the-art methods and future perspectives for personal recognition based on electroencephalogram signals. *Biometrics* 4, 3 (Sept. 2015), 179–190. DOI : <https://doi.org/10.1049/iet-bmt.2014.0040>
- [3] Mohamed Abou-Zleikha, Zheng-Hua Tan, Mads Graesboll Christensen, and Soren Holdt Jensen. 2015. A discriminative approach for speaker selection in speaker de-identification systems. In *European Signal Processing Conference (EUSIPCO)*. IEEE, 2102–2106. DOI : <https://doi.org/10.1109/eusipco.2015.7362755>
- [4] Richard A. Abrams, David E. Meyer, and Sylvan Kornblum. 1989. Speed and accuracy of saccadic eye movements: Characteristics of impulse variability in the oculomotor system. *J. Exp. Psychol. Hum. Percept. Perform.* 15, 3 (1989), 529. DOI : <https://doi.org/10.1037/0096-1523.15.3.529>
- [5] Christopher Ackad, Andrew Clayphan, Roberto Martinez Maldonado, and Judy Kay. 2012. Seamless and continuous user identification for interactive tabletops using personal device handshaking and body tracking. In *Extended Abstracts on Human Factors in Computing Systems*. ACM, 1775–1780. DOI : <https://doi.org/10.1145/2212776.2223708>
- [6] Devon Adams, Alseny Bah, Catherine Barwulor, Nureli Musaby, Kadeem Pitkin, and Elissa M. Redmiles. 2018. Ethics emerging: The story of privacy and security perceptions in virtual reality. In *Symposium on Usable Privacy and Security*. USENIX, 427–442.
- [7] Ayush Agarwal, Amitabh Swain, and S. R. Mahadeva Prasanna. 2022. Speaker anonymization for machines using sinusoidal model. In *IEEE International Conference on Signal Processing and Communications (SPCOM)*. 1–5. DOI : <https://doi.org/10.1109/SPCOM55316.2022.9840792>
- [8] Prachi Agrawal and P. J. Narayanan. 2011. Person de-identification in videos. *Trans. Circ. Syst. Video Technol.* 21, 3 (Mar. 2011), 299–310. DOI : <https://doi.org/10.1109/tcsvt.2011.2105551>
- [9] Hafiz Shehbaz Ali, Fakhar ul Hassan, Siddique Latif, Habib Ullah Manzoor, and Junaid Qadir. 2021. Privacy enhanced speech emotion communication using deep learning aided edge computing. In *International Conference on Communications Workshops*. IEEE, 1–5. DOI : <https://doi.org/10.1109/ICCWorkshops50388.2021.9473669>
- [10] Abdulaziz Almeahmadi and Khalil El-Khatib. 2013. The state of the art in electroencephalogram and access control. In *Conference on Communications and Information Technology (ICCIT)*. IEEE, 49–54. DOI : <https://doi.org/10.1109/iccitechnology.2013.6579521>
- [11] Ranya Aloufi, Hamed Haddadi, and David Boyle. 2020. Privacy-preserving voice analysis via disentangled representations. In *Conference on Cloud Computing Security Workshop*. ACM, 1–14. DOI : <https://doi.org/10.1145/3411495.3421355>
- [12] Arwa Al Sultan and Kevin Warwick. 2013. Keystroke dynamics authentication: A survey of free-text methods. *Int. J. Comput. Sci. Issues* 10, 4 (2013), 1.
- [13] Abdulaziz Alzubaidi and Jugal Kalita. 2016. Authentication of smartphone users using behavioral biometrics. *Commun. Surv. Tutor.* 18, 3 (2016), 1998–2026. DOI : <https://doi.org/10.1109/comst.2016.2537748>
- [14] Miguel E. Andrés, Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2013. Geoindistinguishability: Differential privacy for location-based systems. In *ACM Conference on Computer and Communications Security (CCS)*. 901–914.
- [15] Patricia Arias-Cabarcos, Thilo Habrich, Karen Becker, Christian Becker, and Thorsten Strufe. 2021. Inexpensive brainwave authentication: New techniques and insights on user acceptance. In *30th USENIX Security Symposium (USENIX Security)*. 55–72.
- [16] A. Terry Bahill, Michael R. Clark, and Lawrence Stark. 1975. The main sequence, a tool for studying human eye movements. *Math. Biosci.* 24, 3-4 (Jan. 1975), 191–204. DOI : [https://doi.org/10.1016/0025-5564\(75\)90075-9](https://doi.org/10.1016/0025-5564(75)90075-9)

- [17] Fahimeh Bahmaninezhad, Chunlei Zhang, and John Hansen. 2018. Convolutional neural network based speaker de-identification. In *Speaker and Language Recognition Workshop*. ISCA, 255–260. DOI : <https://doi.org/10.21437/odyssey.2018-36>
- [18] Dustin Bales, Pablo A. Tarazaga, Mary Kasarda, Dhruv Batra, A. G. Woolard, J. D. Poston, and V. V. N. S. Malladi. 2016. Gender classification of walkers via underfloor accelerometer measurements. *Internet Things J.* 3, 6 (Dec. 2016), 1259–1266. DOI : <https://doi.org/10.1109/jiot.2016.2582723>
- [19] Salil Partha Banerjee and Damon Woodard. 2012. Biometric authentication and identification using keystroke dynamics: A survey. *J. Pattern Recog. Res.* 7, 1 (2012), 116–139. DOI : <https://doi.org/10.13176/11.427>
- [20] Sebastian Peter Bayerl, Alexander Wolff von Gudenberg, Florian Hönig, Elmar Noeth, and Korbinian Riedhammer. 2022. KSoF: The Kassel State of Fluency dataset—A therapy centered dataset of stuttering. In *Language Resources and Evaluation Conference*. European Language Resources Association, 1780–1787.
- [21] Henri Begleiter. 1999. EEG Database Data Set. Retrieved from <https://archive.ics.uci.edu/ml/datasets/EEG+Database>
- [22] BehaviorSec. 2019. Continuous authentication through behavioral biometrics. Retrieved from <https://www.behaviosec.com>
- [23] Zineb Bennis and Pierre-Antoine Gourraud. 2021. Application of a novel anonymization method for electrocardiogram data. In *International Conference on Arab Women in Computing*. ACM, 1–5. DOI : <https://doi.org/10.1145/3485557.3485581>
- [24] Shlomo Berkovsky, Ronnie Taib, Irena Koprinska, Eileen Wang, Yucheng Zeng, Jingjie Li, and Sabina Kleitman. 2019. Detecting personality traits using eye-tracking data. In *Conference on Human Factors in Computing Systems (CHI)*. ACM, 1–12. DOI : <https://doi.org/10.1145/3290605.3300451>
- [25] David Bethge, Philipp Hallgarten, Tobias Grosse-Puppenthal, Mohamed Kari, Ralf Mikut, Albrecht Schmidt, and Ozan Ozdenizci. 2022. Domain-invariant representation learning from EEG with private encoders. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1236–1240. DOI : <https://doi.org/10.1109/ICASSP43922.2022.9747398>
- [26] G. Bienvenu and L. Kopp. 1980. Adaptivity to background noise spatial coherence for high resolution passive methods. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 307–310. DOI : <https://doi.org/10.1109/icassp.1980.1171029>
- [27] Stefan Billeb, Christian Rathgeb, Herbert Reininger, Klaus Kasper, and Christoph Busch. 2015. Biometric template protection for speaker recognition based on universal background models. *Biometrics* 4, 2 (June 2015), 116–126. DOI : <https://doi.org/10.1049/iet-bmt.2014.0031>
- [28] Scott Blunsden and R. B. Fisher. 2010. The BEHAVE video dataset: Ground truthed video for multi-person behavior classification. *Ann. BMVA* 4, 1-12 (2010), 4.
- [29] Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. 2012. The quest to replace passwords: A framework for comparative evaluation of web authentication schemes. In *Symposium on Security and Privacy*. IEEE, 553–567. DOI : <https://doi.org/10.1109/sp.2012.44>
- [30] Joseph Bonneau, Cormac Herley, Paul C. van Oorschot, and Frank Stajano. 2015. Passwords and the evolution of imperfect authentication. *Commun. ACM* 58, 7 (June 2015), 78–87. DOI : <https://doi.org/10.1145/2699390>
- [31] Zillah Boraston and Sarah-Jayne Blakemore. 2007. The application of eye-tracking technology in the study of autism. *Physiol. J.* 581, 3 (June 2007), 893–898. DOI : <https://doi.org/10.1113/jphysiol.2007.133587>
- [32] R. Boussetjot, D. Kreiseler, and A. Schnabel. 1995. Nutzung der EKG-signaldatenbank CARDIODAT der PTB über das internet. *Biomediz. Technik* 40, 1 (1995).
- [33] Efe Bozkir, Onur Günlü, Wolfgang Fuhl, Rafael F. Schaefer, and Enkelejda Kasneci. 2021. Differential privacy for eye tracking with temporal correlations. *PLoS ONE* 16, 8 (August 2021), e0255979. DOI : <https://doi.org/10.1371/journal.pone.0255979>
- [34] Attaullah Buriro, Zahid Akhtar, Bruno Crispo, and Filippo Del Frari. 2016. Age, gender and operating-hand estimation on smart mobile devices. In *International Conference of the Biometrics Special Interest Group (BIOSIG)*. IEEE, 1–5. DOI : <https://doi.org/10.1109/biosig.2016.7736910>
- [35] Tom Bäckström, Okko Räsänen, Abraham Zewoudie, and Pablo Pérez Zarazaga. 2021. Introduction to speech processing. WebPage. Retrieved from <https://wiki.aalto.fi/display/ITSP/>
- [36] W. M. Campbell, D. E. Sturim, and D. A. Reynolds. 2006. Support vector machines using GMM supervectors for speaker verification. *IEEE Signal Process. Lett.* 13, 5 (May 2006), 308–311. DOI : <https://doi.org/10.1109/lsp.2006.870086>
- [37] Joseph Campbell and Alan Higgins. 1994. YOHO speaker verification corpus. In *Linguistic Data Consortium*. DOI : <https://doi.org/10.35111/3WC3-N668>
- [38] Emmanuel J. Candes, Justin Romberg, and Terence Tao. 2006. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Trans. Inf. Theor.* 52, 2 (Feb. 2006), 489–509. DOI : <https://doi.org/10.1109/tit.2005.862083>

- [39] Anne M. P. Canuto, Fernando Pintro, and Michael C. Fairhurst. 2014. An effective template protection method for face and voice cancellable identification. *Int. J. Hybrid Intell. Syst.* 11, 3 (2014), 157–166. DOI : <https://doi.org/10.3233/HIS-140192>
- [40] Hyung-Pil Chang, In-Chul Yoo, Changhyeon Jeong, and Dongsuk Yook. 2022. Zero-shot unseen speaker anonymization via voice conversion. *IEEE Access* 10 (2022), 130190–130199. DOI : <https://doi.org/10.1109/ACCESS.2022.3227963>
- [41] Jagmohan Chauhan, Yining Hu, Suranga Seneviratne, Archan Misra, Aruna Seneviratne, and Youngki Lee. 2017. BreathPrint: Breathing acoustics-based user authentication. In *Conference on Mobile Systems, Applications, and Services*. ACM, 278–291. DOI : <https://doi.org/10.1145/3081333.3081355>
- [42] Jagmohan Chauhan, Suranga Seneviratne, Yining Hu, Archan Misra, Aruna Seneviratne, and Youngki Lee. 2018. Breathing-based authentication on resource-constrained IoT devices using recurrent neural networks. *Computer* 51, 5 (May 2018), 60–67. DOI : <https://doi.org/10.1109/mc.2018.2381119>
- [43] Meng Chen, Li Lu, Junhao Wang, Jiadi Yu, Yingying Chen, Zhibo Wang, Zhongjie Ba, Feng Lin, and Kui Ren. 2023. VoiceCloak: Adversarial example enabled voice de-identification with balanced privacy and utility. *Proc. ACM Interact. Mob. Wear. Ubiqu. Technol.* 7, 2, Article 48 (June 2023), 21 pages. DOI : <https://doi.org/10.1145/3596266>
- [44] Peng-Tzu Chen, Shun-Chi Wu, and Jui-Hsuan Hsieh. 2017. A cancelable biometric scheme based on multi-lead ECGs. In *Conference of Engineering in Medicine and Biology Society (EMBC)*. IEEE, 3497–3500. DOI : <https://doi.org/10.1109/embc.2017.8037610>
- [45] Ming Cheng, Xingjian Diao, Shitong Cheng, and Wenjun Liu. 2024. SAIC: Integration of speech anonymization and identity classification. In *AI for Health Equity and Fairness: Leveraging AI to Address Social Determinants of Health*. Springer, 295–306.
- [46] Ching-Yao Chou, En-Jui Chang, Huai-Ting Li, and An-Yeu Wu. 2018. Low-complexity privacy-preserving compressive analysis using subspace-based dictionary for ECG telemonitoring system. *IEEE Trans. Biomed. Circ. Syst.* 12, 4 (Aug. 2018), 801–811. DOI : <https://doi.org/10.1109/tbcas.2018.2828031>
- [47] Oubaida Chouchane, Michele Panariello, Chiara Galdi, Massimiliano Todisco, and Nicholas Evans. 2023. Fairness and privacy in voice biometrics: A study of gender influences using wav2vec 2.0. In *International Conference of the Biometrics Special Interest Group (BIOSIG)*. 1–7. DOI : <https://doi.org/10.1109/BIOSIG58226.2023.10345975>
- [48] V. Ciriani, S. De Capitani di Vimercati, S. Foresti, and P. Samarati. 2008. k-anonymous data mining: A survey. In *Privacy-Preserving Data Mining*. Springer US, 105–136. DOI : https://doi.org/10.1007/978-0-387-70992-5_5
- [49] Alice Cohen-Hadria, Mark Cartwright, Brian McFee, and Juan Pablo Bello. 2019. Voice anonymization in urban sound recordings. In *Workshop on Machine Learning for Signal Processing*. IEEE, 1–6. DOI : <https://doi.org/10.1109/mlsp.2019.8918913>
- [50] Cristina Conati, Christina Merten, Saleema Amershi, and Kasia Muldner. 2007. Using eye-tracking data for high-level user modeling in adaptive interfaces. In *AAAI Conference on Artificial Intelligence (AAAI)*. 1614–1617.
- [51] Emiliano De Cristofaro. 2021. A critical overview of privacy in machine learning. *IEEE Secur. Privac.* 19, 4 (July 2021), 19–27. DOI : <https://doi.org/10.1109/msec.2021.3076443>
- [52] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. 2021. The EPIC-KITCHENS dataset: collection, challenges and baselines. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 11 (2021), 4125–4141. DOI : <https://doi.org/10.1109/TPAMI.2020.2991965>
- [53] Antitza Dantcheva, Petros Elia, and Arun Ross. 2016. What else does your biometric data reveal? A survey on soft biometrics. *IEEE Trans. Inf. Forens. Secur.* 11, 3 (Mar. 2016), 441–467. DOI : <https://doi.org/10.1109/tifs.2015.2480381>
- [54] Badhan Chandra Das, M. Hadi Amini, and Yanzhao Wu. 2025. Security and privacy challenges of large language models: A survey. *ACM Comput. Surv.* 57, 6, Article 152 (Feb. 2025), 39 pages. DOI : <https://doi.org/10.1145/3712001>
- [55] Brendan David-John, Kevin Butler, and Eakta Jain. 2022. For your eyes only: Privacy-preserving eye-tracking datasets. In *Symposium on Eye Tracking Research and Applications*. ACM, 1–6. DOI : <https://doi.org/10.1145/3517031.3529618>
- [56] Brendan David-John, Kevin Butler, and Eakta Jain. 2023. Privacy-preserving datasets of eye-tracking samples with applications in XR. *IEEE Trans. Visualiz. Comput. Graph.* 29, 5 (2023), 2774–2784. DOI : <https://doi.org/10.1109/TVCG.2023.3247048>
- [57] Brendan David-John, Diane Hofelt, Kevin Butler, and Eakta Jain. 2021. A privacy-preserving approach to streaming eye-tracking data. *IEEE Transactions on Visualization and Computer Graphics* 27, 5 (May 2021), 2555–2565. DOI : <https://doi.org/10.1109/tvcg.2021.3067787>
- [58] Maria Cecilia Teixeira de Carvalho Bruno, Maria Aparecida Constantino Vilela, and Carlos Alberto B. Mendes de Oliveira. 2013. Study on dermatoses and their prevalence in groups of confirmed alcoholic individuals in comparison to a non-alcoholic group of individuals. *Anais Brasil. Dermatol.* 88, 3 (June 2013), 368–375. DOI : <https://doi.org/10.1590/abd1806-4841.20131829>
- [59] Ana Lgia Silva de Lima, Luc J. W. Evers, Tim Hahn, Lauren Bataille, Jamie L. Hamilton, Max A. Little, Yasuyuki Okuma, Bastiaan R. Bloem, and Marjan J. Faber. 2017. Freezing of gait and fall detection in Parkinson’s disease using

- wearable sensors: A systematic review. *J. Neurol.* 264, 8 (Mar. 2017), 1642–1654. DOI : <https://doi.org/10.1007/s00415-017-8424-0>
- [60] Wheidima Carneiro De Melo, Eric Granger, and Abdenour Hadid. 2019. Depression detection based on deep distribution learning. In *IEEE International Conference on Image Processing (ICIP)*. IEEE, 4544–4548.
- [61] Yves-Alexandre de Montjoye, César A. Hidalgo, Michel Verleysen, and Vincent D. Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scient. Rep.* 3, 1 (Mar. 2013), 1376. DOI : <https://doi.org/10.1038/srep01376>
- [62] Essam Debie, Nour Moustafa, and Monica T. Whitty. 2020. A privacy-preserving generative adversarial network method for securing EEG brain signals. In *International Joint Conference on Neural Networks*. IEEE, 1–8. DOI : <https://doi.org/10.1109/IJCNN48605.2020.9206683>
- [63] Noëlie Debs, Théo Jourdan, Ali Moukadem, Antoine Boutet, and Carole Frindel. 2021. Motion sensor data anonymization by time-frequency filtering. In *28th European Signal Processing Conference (EUSIPCO)*. 1707–1711. DOI : <https://doi.org/10.23919/Eusipco47968.2020.9287683>
- [64] Najim Dehak, Patrick J. Kenny, Réda Dehak, Pierre Dumouchel, and Pierre Ouellet. 2011. Front-end factor analysis for speaker verification. *Trans. Audio, Speech, Lang. Process.* 19, 4 (May 2011), 788–798. DOI : <https://doi.org/10.1109/tasl.2010.2064307>
- [65] Jiangyi Deng, Fei Teng, Yanjiao Chen, Xiaofu Chen, Zhaohui Wang, and Wenyuan Xu. 2023. V-Cloak: Intelligibility-, naturalness- & timbre-preserving real-time voice anonymization. In *32nd USENIX Security Symposium (USENIX Security)*. 5181–5198.
- [66] Li Deng. 2012. The MNIST database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Process. Mag.* 29, 6 (Nov. 2012), 141–142. DOI : <https://doi.org/10.1109/msp.2012.2211477>
- [67] Joy Derwenskus, Janet C. Rucker, Alessandro Serra, John S. Stahl, Deborah L. Downey, Nancy L. Adams, and R. John Leigh. 2005. Abnormal eye movements predict disability in MS: Two-year follow-up. *Ann. New York Acad. Sci.* 1039, 1 (Apr. 2005), 521–523. DOI : <https://doi.org/10.1196/annals.1325.058>
- [68] Clemens Deuser, Steffen Passmann, and Thorsten Strufe. 2020. Browsing unicity: On the limits of anonymizing web tracking data. In *Symposium on Security and Privacy*. IEEE, 279–292. DOI : <https://doi.org/10.1109/sp40000.2020.00018>
- [69] Apiwat Dittthapron, Emmanuel O. Agu, and Adam C. Lammert. 2021. Privacy-preserving deep speaker separation for smartphone-based passive speech assessment. *IEEE Open J. Eng. Med. Biol.* 2 (2021), 304–313. DOI : <https://doi.org/10.1109/OJEMB.2021.3063994>
- [70] Hamza Djelouat, Xiaojun Zhai, Mohamed Al Disi, Abbes Amira, and Faycal Bensaali. 2018. System-on-chip solution for patients biometric: A compressive sensing-based approach. *IEEE Sensors J.* 18, 23 (Dec. 2018), 9629–9639. DOI : <https://doi.org/10.1109/jsen.2018.2871411>
- [71] Isha Dua, Thrupthi Ann John, Riya Gupta, and C. V. Jawahar. 2020. DGAZE: Driver gaze mapping on road. In *Conference on Intelligent Robots and Systems*.
- [72] Andrew T. Duchowski. 2017. *Eye Tracking Methodology*. Springer International Publishing. DOI : <https://doi.org/10.1007/978-3-319-57883-5>
- [73] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2017. Calibrating noise to sensitivity in private data analysis. *J. Privac. Confid.* 7, 3 (May 2017), 17–51. DOI : <https://doi.org/10.29012/jpc.v7i3.405>
- [74] Cynthia Dwork and Aaron Roth. 2013. The algorithmic foundations of differential privacy. *Found. Trends Theoret. Comput. Sci.* 9, 3-4 (2013), 211–407. DOI : <https://doi.org/10.1561/04000000042>
- [75] Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. 2017. Exposed! A survey of attacks on private data. *Ann. Rev. Stat. Applic.* 4, 1 (Mar. 2017), 61–84. DOI : <https://doi.org/10.1146/annurev-statistics-060116-054123>
- [76] Simon Eberz, Giulio Lovisotto, Andrea Patane, Marta Kwiatkowska, Vincent Lenders, and Ivan Martinovic. 2018. When your fitness tracker betrays you: Quantifying the predictability of biometric features across contexts. In *Symposium on Security and Privacy*. IEEE, 889–905. DOI : <https://doi.org/10.1109/sp.2018.00053>
- [77] Sarah Ebling, Necati Camgoz, Penny Braem, Katja Tissi, Sandra Sidler-Miserez, Stephanie Stoll, Simon Hadfield, Tobias Haug, Richard Bowden, Sandrine Tornay et al. 2018. SMILE Swiss German sign language dataset. In *International Conference on Language Resources and Evaluation*.
- [78] Khaled El Emam, Elizabeth Jonker, Luk Arbuckle, and Bradley Malin. 2011. A systematic review of re-identification attacks on health data. *PLoS One* 6, 12 (Dec. 2011), e28071. DOI : <https://doi.org/10.1371/journal.pone.0028071>
- [79] Fatih Ertam. 2019. An effective gender recognition approach using voice data via deeper LSTM networks. *Appl. Acoust.* 156 (Dec. 2019), 351–358. DOI : <https://doi.org/10.1016/j.apacoust.2019.07.033>
- [80] Ulrich Ettinger, Veena Kumari, Xavier A. Chitnis, Philip J. Corr, Trevor J. Crawford, Dominic G. Fannon, Séamus O’Ceallaigh, Alex L. Sumich, Victor C. Doku, and Tonmoy Sharma. 2004. Volumetric neural correlates of antisaccade eye movements in first-episode psychosis. *Am. J. Psychiat.* 161, 10 (Oct. 2004), 1918–1921. DOI : <https://doi.org/10.1176/ajp.161.10.1918>
- [81] Jiahao Fan and Xiaogang Hu. 2023. Privacy-preserving motor intent classification via feature disentanglement. In *11th International IEEE/EMBS Conference on Neural Engineering (NER)*. 1–4. DOI : <https://doi.org/10.1109/NER52421.2023.10123842>

- [82] Fuming Fang, Xin Wang, Junichi Yamagishi, Isao Echizen, Massimiliano Todisco, Nicholas Evans, and Jean-Francois Bonastre. 2019. Speaker anonymization using x-vector and neural waveform models. In *Speech Synthesis Workshop*. DOI : <https://doi.org/10.21437/ssw.2019-28>
- [83] Marcos Faundez-Zanuy, Enric Sesa-Nogueras, and Stefano Marinuzzi. 2015. Speaker identification experiments under gender de-identification. In *Carnahan Conference on Security Technology*. IEEE, 1–6. DOI : <https://doi.org/10.1109/ccst.2015.7389702>
- [84] J. Fierrez, J. Galbally, J. Ortega-Garcia, M. R. Freire, F. Alonso-Fernandez, D. Ramos, D. T. Toledano, J. Gonzalez-Rodriguez, J. A. Siguenza, J. Garrido-Salas et al. 2009. BiosecuID: A multimodal biometric database. *Pattern Anal. Applic.* 13, 2 (Feb. 2009), 235–246. DOI : <https://doi.org/10.1007/s10044-009-0151-4>
- [85] Lucas Silva Figueiredo, Benjamin Livshits, David Molnar, and Margus Veanes. 2016. Prepose: Privacy, security, and reliability for gesture-based programming. In *Symposium on Security and Privacy*. IEEE, 122–137. DOI : <https://doi.org/10.1109/sp.2016.16>
- [86] Wolfgang Fuhl, Efe Bozkir, and Enkelejd Kasneci. 2021. Reinforcement learning for the privacy preservation and manipulation of eye tracking data. In *International Conference on Artificial Neural Networks*. Springer, 595–607.
- [87] Bence Galai and Csaba Benedek. 2015. Feature selection for lidar-based gait recognition. In *Workshop on Computational Intelligence for Multimedia Understanding*. IEEE, 1–5. DOI : <https://doi.org/10.1109/iwcm.2015.7347076>
- [88] Stephan J. Garbin, Yiru Shen, Immo Schuetz, Robert Cavin, Gregory Hughes, and Sachin S. Talathi. 2019. *OpenEDS: Open Eye Dataset*. arXiv. DOI : <https://doi.org/10.48550/ARXIV.1905.03702>
- [89] Ana Garcia-Blanco, Ladislao Salmerón, Manuel Perea, and Lorenzo Livianos. 2014. Attentional biases toward emotional images in the different episodes of bipolar disorder: An eye-tracking study. *Psychiat. Res.* 215, 3 (Mar. 2014), 628–633. DOI : <https://doi.org/10.1016/j.psychres.2013.12.039>
- [90] Giuseppe Garofalo, Tim Van hamme, Davy Preuveneers, and Wouter Joosen. 2020. A siamese adversarial anonymizer for data minimization in biometric applications. In *European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 334–343. DOI : <https://doi.org/10.1109/EuroSPW51379.2020.00052>
- [91] J. Garofolo, Lori Lamel, W. Fisher, Jonathan Fiscus, D. Pallett, N. Dahlgren, and V. Zue. 1992. TIMIT acoustic-phonetic continuous speech corpus. In *Linguistic Data Consortium*.
- [92] Ünal Ege Gaznepoglu and Nils Peters. 2023. Deep learning-based F0 synthesis for speaker anonymization. In *31st European Signal Processing Conference (EUSIPCO)*. 291–295. DOI : <https://doi.org/10.23919/EUSIPCO58844.2023.10290038>
- [93] Romain Giot, Mohamad El-Abed, and Christophe Rosenberger. 2009. GREYC keystroke: A benchmark for keystroke dynamics biometric systems. In *Biometrics: Theory, Applications, and Systems*. IEEE, 1–6. DOI : <https://doi.org/10.1109/btas.2009.5339051>
- [94] Romain Giot, Mohamad El-Abed, and Christophe Rosenberger. 2012. Web-based benchmark for keystroke dynamics biometric systems: A statistical analysis. In *Intelligent Information Hiding and Multimedia Signal Processing*. IEEE, 11–15. DOI : <https://doi.org/10.1109/iih-msp.2012.10>
- [95] Abenezer Golda, Kidus Mekonen, Amit Pandey, Anushka Singh, Vikas Hassija, Vinay Chamola, and Biplab Sikdar. 2024. Privacy and security concerns in generative AI: A comprehensive survey. *IEEE Access* 12 (2024), 48126–48144. DOI : <https://doi.org/10.1109/ACCESS.2024.3381611>
- [96] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (Oct. 2020), 139–144. DOI : <https://doi.org/10.1145/3422622>
- [97] Yuuki Goubaru, Yasushi Yamazaki, Takeru Miyazaki, and Tetsushi Ohki. 2014. A consideration on a common template-based biometric cryptosystem using on-line signatures. In *Global Conference on Consumer Electronics*. IEEE, 131–135. DOI : <https://doi.org/10.1109/gcce.2014.7031229>
- [98] Erin Griffiths, Salah Assana, and Kamin Whitehouse. 2018. Privacy-preserving image processing with binocular thermal cameras. *Interact., Mob., Wear. Ubiqu. Technol.* 1, 4 (Jan. 2018), 1–25. DOI : <https://doi.org/10.1145/3161198>
- [99] Qiong Gui, Maria V. Ruiz-Blondet, Sarah Laszlo, and Zhanpeng Jin. 2019. A survey on brain biometrics. *Comput. Surv.* 51, 6 (Feb. 2019), 1–38. DOI : <https://doi.org/10.1145/3230632>
- [100] Priyanka Gupta, Gauri P. Prajapati, Shrishti Singh, Madhu R. Kamble, and Hemant A. Patil. 2020. Design of voice privacy system using linear prediction. In *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference, APSIPA*, IEEE, 543–549. Retrieved from <https://ieeexplore.ieee.org/document/9306379>
- [101] Lindsay F. Haas. 2003. Hans Berger (1873–1941), Richard Caton (1842–1926), and electroencephalography. *J. Neurol. Neurosurg. Psychiat.* 74, 1 (Jan. 2003), 9–9. DOI : <https://doi.org/10.1136/jnnp.74.1.9>
- [102] Agrya Halder, Pratik Chattopadhyay, and Sathish Kumar. 2023. Gait transformation network for gait de-identification with pose preservation. *Signal, Image Video Process.* 17, 5 (2023), 1753–1761.
- [103] Jihun Hamm. 2017. Enhancing utility and privacy with noisy minimax filters. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 6389–6393. DOI : <https://doi.org/10.1109/icassp.2017.7953386>

- [104] Yaowei Han, Sheng Li, Yang Cao, Qiang Ma, and Masatoshi Yoshikawa. 2020. Voice-indistinguishability: Protecting voiceprint in privacy-preserving speech data release. In *Conference on Multimedia and Expo (ICME)*. IEEE, 1–6. DOI: <https://doi.org/10.1109/ICME46284.2020.9102875>
- [105] Simon Hanisch, Evelyn Muschter, Admantini Hatzipanayioti, Shu-Chen Li, and Thorsten Strufe. 2023. Understanding person identification through gait. *Proc. Privac. Enhanc. Technol.* 1 (2023), 177–189.
- [106] Katarzyna Harezlak and Pawel Kasprowski. 2018. Application of eye tracking in medicine: A survey, research issues and challenges. *Comput. Med. Imag. Graph.* 65 (Apr. 2018), 176–190. DOI: <https://doi.org/10.1016/j.compmedimag.2017.04.006>
- [107] Kei Hashimoto, Junichi Yamagishi, and Isao Echizen. 2016. Privacy-preserving sound to degrade automatic speaker verification performance. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 5500–5504. DOI: <https://doi.org/10.1109/icassp.2016.7472729>
- [108] Jane Henriksen-Bulmer and Sheridan Jeary. 2016. Re-identification attacks—A systematic literature review. *Int. J. Inf. Manag.* 36, 6 (Dec. 2016), 1184–1192. DOI: <https://doi.org/10.1016/j.ijinfomgt.2016.08.002>
- [109] Eckhard H. Hess and James M. Polt. 1960. Pupil size as related to interest value of visual stimuli. *Science* 132, 3423 (Aug. 1960), 349–350. DOI: <https://doi.org/10.1126/science.132.3423.349>
- [110] Jan Hintz, Sebastian Bayerl, Yamini Sinha, Suhita Ghosh, Martha Schubert, Sebastian Stober, Korbinian Riedhammer, and Ingo Siegert. 2023. Anonymization of stuttered speech—Removing speaker information while preserving the utterance. In *3rd Symposium on Security and Privacy in Speech Communication*. 41–45. DOI: <https://doi.org/10.21437/SPSC.2023-7>
- [111] HIPAA Compliance Assistance. 2003. *Summary of the HIPAA Privacy Rule*. Publication of the US Dept. of Health.
- [112] Yuki Hirose, Kazuaki Nakamura, Naoko Nitta, and Noboru Babaguchi. 2019. Anonymization of gait silhouette video by perturbing its phase and shape components. In *Asia-Pacific Signal and Information Processing Association Annual Summit*. IEEE, 1679–1685. DOI: <https://doi.org/10.1109/apsipaasc47483.2019.9023196>
- [113] Thang Hoang, Deokjai Choi, and Thuc Nguyen. 2015. Gait authentication on mobile phone using biometric cryptosystem and fuzzy commitment scheme. *Int. J. Inf. Secur.* 14, 6 (Jan. 2015), 549–560. DOI: <https://doi.org/10.1007/s10207-015-0273-1>
- [114] Ulrich Hoffmann, Jean-Marc Vesin, Touradj Ebrahimi, and Karin Diserens. 2008. An efficient P300-based brain–computer interface for disabled subjects. *J. Neurosci. Meth.* 167, 1 (2008), 115–125. DOI: <https://doi.org/10.1016/j.jneumeth.2007.03.005>
- [115] Giles Hogben. 2010. ENISA briefing: Behavioural biometrics. *Computat. Intell.* (2010).
- [116] Philip S. Holzman, Leonard R. Proctor, and Dominic W. Hughes. 1973. Eye-tracking patterns in schizophrenia. *Science* 181, 4095 (July 1973), 179–181. DOI: <https://doi.org/10.1126/science.181.4095.179>
- [117] Pei-Lun Hong, Jyun-Ya Hsiao, Chi-Hsun Chung, Yao-Min Feng, and Shun-Chi Wu. 2019. ECG biometric recognition: Template-free approaches based on deep learning. In *Annual International Conference of Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2633–2636. DOI: <https://doi.org/10.1109/embc.2019.8856916>
- [118] Syed Monowar Hossain, Amin Ahsan Ali, Md. Mahbubur Rahman, Emre Ertine David Epstein, Ashley Kennedy, Kenzie Preston, Annie Umbricht, Yixin Chen, and Santosh Kumar. 2014. Identifying drug (cocaine) intake events from acute physiological response in the presence of free-living physical activity. In *International Symposium on Information Processing in Sensor Networks*. IEEE, 71–82. DOI: <https://doi.org/10.1109/ipsn.2014.6846742>
- [119] Miao Hu, Zhenxiao Luo, Yipeng Zhou, Xuezheng Liu, and Di Wu. 2022. Otus: A gaze model-based privacy control framework for eye tracking applications. In *Conference on Computer Communications (INFOCOM)*. IEEE, 560–569. DOI: <https://doi.org/10.1109/INFOCOM48880.2022.9796665>
- [120] Zhiming Hu, Andreas Bulling, Sheng Li, and Guoping Wang. 2023. EHTask: Recognizing user tasks From eye and head movements in immersive virtual reality. *IEEE Transactions on Visualization and Computer Graphics* 29, 4 (April 2023), 1992–2004. DOI: <https://doi.org/10.1109/tvcg.2021.3138902>
- [121] Pei Huang, Linke Guo, Ming Li, and Yuguang Fang. 2019. Practical privacy-preserving ECG-based authentication for IoT-based healthcare. *IEEE Internet Things J.* 6, 5 (Oct. 2019), 9200–9210. DOI: <https://doi.org/10.1109/jiot.2019.2929087>
- [122] Anco Hundepool, Josep Domingo-Ferrer, Luisa Franconi, Sarah Giessing, Eric S. Nordholt, Keith Spicer, and Peter-Paul de Wolf. 2012. *Statistical Disclosure Control*. Wiley.
- [123] J. Thomas Hutton, J. A. Nagel, and Ruth B. Loewenson. 1984. Eye tracking dysfunction in Alzheimer-type dementia. *Neurol.* 34, 1 (Jan. 1984), 99–99. DOI: <https://doi.org/10.1212/wnl.34.1.99>
- [124] Michiko Inoue, Masashi Nishiyama, and Yoshio Iwai. 2020. Gender classification using the gaze distributions of observers on privacy-protected training images. In *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*. SCITEPRESS, 149–156. DOI: <https://doi.org/10.5220/0008876101490156>
- [125] M. Ivasic-Kos, A. Iosifidis, A. Tefas, and I. Pitas. 2014. Person de-identification in activity videos. In *Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. IEEE, 1294–1299. DOI: <https://doi.org/10.1109/mipro.2014.6859767>

- [126] Salar Jafarlou, Amir M. Rahmani, Nikil Dutt, and Sanaz Rahimi Mousavi. 2022. ECG biosignal deidentification using conditional generative adversarial networks. In *44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. 1366–1370. DOI : <https://doi.org/10.1109/EMBC48229.2022.9872015>
- [127] Jinhyeok Jang, Dohyung Kim, Cheonshu Park, Minsu Jang, Jaeyeon Lee, and Jaehong Kim. 2020. ETRI-Activity3D: A large-scale RGB-D dataset for robots to recognize daily activities of the elderly. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE Press, 10990–10997. DOI : <https://doi.org/10.1109/IROS45743.2020.9341160>
- [128] J. Jankovic. 2008. Parkinson’s disease: Clinical features and diagnosis. *J. Neurol. Neurosurg. Psychiat.* 79, 4 (2008), 368–376. DOI : <https://doi.org/10.1136/jnnp.2007.131045>
- [129] Qin Jin, Arthur R. Toth, Tanja Schultz, and Alan W. Black. 2009. Voice convergin: Speaker de-identification by voice transformation. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 3909–3912. DOI : <https://doi.org/10.1109/icassp.2009.4960482>
- [130] I. Joe Louis Paul, S. Sasirekha, S. Uma Maheswari, K. A. M. Ajith, S. M. Arjun, and S. Athesh Kumar. 2019. Eye gaze tracking-based adaptive e-learning for enhancing teaching and learning in virtual classrooms. In *Information and Communication Technology for Competitive Strategies*. Springer, 165–176.
- [131] Théo Jourdan, Antoine Boutet, and Carole Frindel. 2018. Toward privacy in IoT mobile devices for activity recognition. In *International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*. ACM, 155–165. DOI : <https://doi.org/10.1145/3286978.3287009>
- [132] Tadej Justin, Vitomir Struc, Simon Dobrisek, Bostjan Vesnicer, Ivo Ipsic, and France Mihelic. 2015. Speaker de-identification using diphone recognition and speech synthesis. In *Automatic Face and Gesture Recognition*. IEEE, 1–7. DOI : <https://doi.org/10.1109/fg.2015.7285021>
- [133] E. Grace Mary Kanaga, R. Muthu Kumaran, M. Hema, R. Gowri Manohari, and Tina Anu Thomas. 2017. An experimental investigations on classifiers for brain computer interface (BCI) based authentication. In *Conference on Trends in Electronics and Informatics (ICEI)*. IEEE, 1–6. DOI : <https://doi.org/10.1109/icoei.2017.8300873>
- [134] Nader Karamzadeh, Yasaman Ardeshirpour, Matthew Kellman, Fatima Chowdhry, Afroz Anderson, David Chorlian, Edward Wegman, and Amir Gandjbakhche. 2015. Relative brain signature: A population-based feature extraction procedure to identify functional biomarkers in the brain of alcoholics. *Brain Behav.* 5, 7 (May 2015), e00335. DOI : <https://doi.org/10.1002/brb3.335>
- [135] Stamos Katsigiannis and Naeem Ramzan. 2017. DREAMER: A database for emotion recognition through EEG and ECG signals from wireless low-cost off-the-shelf devices. DOI : <https://doi.org/10.1109/JBHI.2017.2688239>
- [136] Christina Katsini, Yasmeen Abdrabou, George E. Raptis, Mohamed Khamis, and Florian Alt. 2020. The role of eye gaze in security and privacy applications: Survey and future HCI research directions. In *Conference on Human Factors in Computing Systems (CHI)*. ACM, 1–21. DOI : <https://doi.org/10.1145/3313831.3376840>
- [137] Georgios Kellaris, Stavros Papadopoulos, Xiaokui Xiao, and Dimitris Papadias. 2014. Differentially private event sequences over infinite streams. *Proc. VLDB Endow.* 7, 12 (2014), 1155–1166.
- [138] Gokce Keskin, Tyler Lee, Cory Stephenson, and Oguz H. Elibol. 2019. *Measuring the Effectiveness of Voice Conversion on Speaker Identification and Automatic Speech Recognition Systems*. arXiv. DOI : <https://doi.org/10.48550/arxiv.1905.12531>
- [139] W. Khalifa, A. Salem, and M. Roushdy. 2012. A survey of EEG based user authentication schemes. In *International Conference on INFormatics and Systems*. 55–60.
- [140] Christopher Kirtley. 2006. *Clinical Gait Analysis: Theory and Practice*. Elsevier Health Sciences.
- [141] Barbara Kitchenham. 2004. *Procedures for Performing Systematic Reviews*. Technical Report TR/SE-0401. Keele University, Keele, UK.
- [142] Sander Koelstra, Christian Muhl, Mohammad Soleymani, Jong-Seok Lee, Ashkan Yazdani, Touradj Ebrahimi, Thierry Pun, Anton Nijholt, and Ioannis Patras. 2012. DEAP: A database for emotion analysis using physiological signals. *Trans. Affect. Comput.* 3, 1 (2012), 18–31. DOI : <https://doi.org/10.1109/T-AFFC.2011.15>
- [143] Kazuhiro Kondo, Tomohiro Komiya, and Shintaro Kashiwada. 2013. Towards gender-dependent babble maskers for speech privacy protection. In *Conference on Intelligent Information Hiding and Multimedia Signal Processing*. IEEE, 275–278. DOI : <https://doi.org/10.1109/iih-msp.2013.77>
- [144] Kazuhiro Kondo and Hiroki Sakurai. 2014. Gender-dependent babble maskers created from multi-speaker speech for speech privacy protection. In *Conference on Intelligent Information Hiding and Multimedia Signal Processing*. IEEE, 251–254. DOI : <https://doi.org/10.1109/iih-msp.2014.69>
- [145] M. Kosinski, D. Stillwell, and T. Graepel. 2013. Private traits and attributes are predictable from digital records of human behavior. *Proc. Nat’l Acad. Sci.* 110, 15 (Mar. 2013), 5802–5805. DOI : <https://doi.org/10.1073/pnas.1218772110>
- [146] Krzysztof Krejtz, Andrew T. Duchowski, Anna Niedzielska, Cezary Biele, and Izabela Krejtz. 2018. Eye tracking cognitive load using pupil diameter and microsaccades with fixed gaze. *PLoS One* 13, 9 (Sept. 2018), e0203629. DOI : <https://doi.org/10.1371/journal.pone.0203629>

- [147] Jacob Leon Kröger, Otto Hans-Martin Lutz, and Florian Müller. 2020. What does your gaze reveal about you? On the privacy implications of eye tracking. In *Privacy and Identity Management. Data for Better Living: AI and Privacy*. Springer International Publishing, 226–241. DOI : https://doi.org/10.1007/978-3-030-42504-3_15
- [148] Craig A. Kuechenmeister, Patrick H. Linton, Thelma V. Mueller, and Hilton B. White. 1977. Eye tracking in relation to age, sex, and illness. *Arch. Gen. Psychiatry* 34, 5 (May 1977), 578–579. DOI : <https://doi.org/10.1001/archpsyc.1977.01770170088008>
- [149] Lamyamba Laishram, Muhammad Shaheryar, Jong Taek Lee, and Soon Ki Jung. 2025. Toward a privacy-preserving face recognition system: A survey of leakages and solutions. *ACM Comput. Surv.* 57, 6, Article 147 (Feb. 2025), 38 pages. DOI : <https://doi.org/10.1145/3673224>
- [150] Anthony Larcher, Kong Aik Lee, Bin Ma, and Haizhou Li. 2012. The RSR2015: Database for text-dependent speaker verification using multiple pass-phrases. In *13th Annual Conference of the International Speech Communication Association (INTERSPEECH)*. 1578–1581.
- [151] Juho Leinonen, Petri Ihanntola, and Arto Hellas. 2017. Preventing keystroke based identification in open data sets. In *Conference on Learning @ Scale*. ACM, 101–109. DOI : <https://doi.org/10.1145/3051457.3051458>
- [152] Deborah L. Levy, Anne B. Sereno, Diane C. Gooding, and Gillian A. O’Driscoll. 2010. Eye tracking dysfunction in schizophrenia: Characterization and pathophysiology. In *Behavioral Neurobiology of Schizophrenia and Its Treatment*. Springer, 311–347. DOI : https://doi.org/10.1007/7854_2010_60
- [153] Jingjie Li, Amrita Roy Chowdhury, Kassem Fawaz, and Younghyun Kim. 2021. Kaledo: Real-time privacy control for eye-tracking systems. In *USENIX Security Conference*. 1793–1810. Retrieved from <https://www.usenix.org/conference/usenixsecurity21/presentation/li-jingjie>
- [154] Yunji Liang, Sagar Samtani, Bin Guo, and Zhiwen Yu. 2020. Behavioral biometrics for continuous authentication in the internet-of-things era: An artificial intelligence perspective. *Internet Things J.* 7, 9 (Sept. 2020), 9128–9143. DOI : <https://doi.org/10.1109/jiot.2020.3004077>
- [155] Jae Lim and A. Oppenheim. 1978. All-pole modeling of degraded speech. *Trans. Audio, Speech, Lang. Process.* 26, 3 (June 1978), 197–210. DOI : <https://doi.org/10.1109/tassp.1978.1163086>
- [156] Ao Liu, Lirong Xia, Andrew Duchowski, Reynold Bailey, Kenneth Holmqvist, and Eakta Jain. 2019. Differential privacy for eye-tracking data. In *Symposium on Eye Tracking Research & Applications*. ACM, 1–10. DOI : <https://doi.org/10.1145/3314111.3319823>
- [157] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C. Kot. 2020. NTU RGB+D 120: A large-scale benchmark for 3D human activity understanding. *IEEE Trans. Pattern Anal. Mach. Intell.* 42, 10 (2020), 2684–2701.
- [158] Xinwen Liu, Huan Wang, Zongjin Li, and Lang Qin. 2021. Deep learning in ECG diagnosis: A review. *Knowl.-based Syst.* 227 (2021), 107187. DOI : <https://doi.org/10.1016/j.knosys.2021.107187>
- [159] Wen-Chih Lo, Ching-Ling Fan, Jean Lee, Chun-Ying Huang, Kuan-Ta Chen, and Cheng-Hsin Hsu. 2017. 360° video viewing dataset in head-mounted virtual reality. In *Multimedia Systems Conference (MMSys)*. ACM, New York, NY, USA, 211–216. DOI : <https://doi.org/10.1145/3083187.3083219>
- [160] Dillon Lohr, Samantha Aziz, Lee Friedman, and Oleg V. Komogortsev. 2023. GazeBaseVR, a large-scale, longitudinal, binocular eye-tracking dataset collected in virtual reality. *Scient. Data* 10, 1 (2023), 177.
- [161] Paula Lopez-Otero, Carmen Magariños, Laura Docio-Fernandez, Eduardo Rodriguez-Banga, Daniel Erro, and Carmen Garcia-Mateo. 2017. Influence of speaker de-identification in depression detection. *Signal Process.* 11, 9 (Dec. 2017), 1023–1030. DOI : <https://doi.org/10.1049/iet-spr.2016.0731>
- [162] Yuanjun Lv, Jixun Yao, Peikun Chen, Hongbin Zhou, Heng Lu, and Lei Xie. 2023. SALT: Distinguishable speaker anonymization through latent space transformation. In *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 1–8.
- [163] Xiaosong Ma, Yubo Song, Zhongwei Wang, Shang Gao, Bin Xiao, and Aiqun Hu. 2021. You can hear but you cannot record: Privacy protection by jamming audio recording. In *International Conference on Communications*. IEEE, 1–6. DOI : <https://doi.org/10.1109/ICC42927.2021.9500456>
- [164] Carmen Magariños, Paula Lopez-Otero, Laura Docio-Fernandez, Eduardo Rodriguez-Banga, Daniel Erro, and Carmen Garcia-Mateo. 2017. Reversible speaker de-identification using pre-trained transformation functions. *Comput. Speech Lang.* 46 (Nov. 2017), 36–52. DOI : <https://doi.org/10.1016/j.csl.2017.05.001>
- [165] Ahmed Mahfouz, Tarek M. Mahmoud, and Ahmed Sharaf Eldin. 2017. A survey on behavioral biometric authentication on smartphones. *J. Inf. Secur. Applic.* 37 (Dec. 2017), 28–37. DOI : <https://doi.org/10.1016/j.jisa.2017.10.002>
- [166] Seedahmed S. Mahmoud. 2016. A generalised wavelet packet-based anonymisation approach for ECG security application. *Secur. Commun. Netw.* 9, 18 (Dec. 2016), 6137–6147. DOI : <https://doi.org/10.1002/sec.1762>
- [167] Emanuele Maiorana, Patrizio Campisi, and Alessandro Neri. 2011. Bioconvolving: Cancelable templates for a multi-biometrics signature recognition system. In *International Systems Conference*. IEEE, 495–500. DOI : <https://doi.org/10.1109/syscon.2011.5929064>

- [168] Anindya Maiti, Oscar Armbruster, Murtuza Jadliwala, and Jibo He. 2016. Smartwatch-based keystroke inference attacks and context-aware protection mechanisms. In *ACM ASIA Conference on Computer and Communications Security (ACM ASIACCS)*. ACM, 795–806. DOI : <https://doi.org/10.1145/2897845.2897905>
- [169] Päivi Majaranta and Andreas Bulling. 2014. Eye tracking and eye-based human–computer interaction. In *Human–computer Interaction*. Springer London, 39–65. DOI : https://doi.org/10.1007/978-1-4471-6392-3_3
- [170] Yasushi Makihara, Hidetoshi Mannami, Akira Tsuji, Md. Altab Hossain, Kazushige Sugiura, Atsushi Mori, and Yasushi Yagi. 2012. The OU-ISIR Gait database comprising the treadmill dataset. *IPSPJ Trans. Comput. Vis. Appl.* 4 (Apr. 2012), 53–62. DOI : <https://doi.org/10.2197/ipsjtcva.4.53>
- [171] Mohammad Malekzadeh, Richard G. Clegg, Andrea Cavallaro, and Hamed Haddadi. 2020. Privacy and utility preserving sensor-data transformations. *Pervasive and Mobile Computing* 63 (March 2020), 101132. DOI : <https://doi.org/10.1016/j.pmcj.2020.101132>
- [172] M. Sabarimalai Manikandan and S. Dandapat. 2008. ECG distortion measures and their effectiveness. In *Emerging Trends in Engineering and Technology*. IEEE, 705–710. DOI : <https://doi.org/10.1109/icetec.2008.248>
- [173] Ivan Martinovic, Doug Davies, Mario Frank, Daniele Perito, Tomas Ros, and Dawn Song. 2012. On the feasibility of side-channel attacks with brain-computer interfaces. In *USENIX Security Conference*. 143–158.
- [174] Richard Matovu and Abdul Serwadda. 2016. Your substance abuse disorder is an open secret! Gleaning sensitive personal information from templates in an EEG-based authentication system. In *International Conference on Biometrics Theory, Applications and Systems*. IEEE, 1–7. DOI : <https://doi.org/10.1109/btas.2016.7791210>
- [175] Richard Matovu, Abdul Serwadda, David Irakiza, and Isaac Griswold-Steiner. 2018. Jekyll and Hyde: On the double-faced nature of smart-phone sensor noise injection. In *International Conference of the Biometrics Special Interest Group (BIOSIG)*. IEEE, 1–6. DOI : <https://doi.org/10.23919/biosig.2018.8553043>
- [176] Gerald Matthews, W. Middleton, Bernard Gilmartin, and Mark A. Bullimore. 1991. Pupillary diameter and cognitive load. *Journal of Psychophysiology* 5 (1991), 265–271.
- [177] Candy Olivia Mawalim, Kasorn Galajit, Jessada Karnjana, Shunsuke Kidani, and Masashi Unoki. 2022. Speaker anonymization by modifying fundamental frequency and x-vector singular value. *Computer Speech & Language* 73 (May 2022), 101326. DOI : <https://doi.org/10.1016/j.csl.2021.101326>
- [178] Candy Olivia Mawalim, Shogo Okada, and Masashi Unoki. 2022. Speaker anonymization by pitch shifting based on time-scale modification. In *2nd Symposium on Security and Privacy in Speech Communication*. 35–42.
- [179] Frank D. McSherry. 2009. Privacy integrated queries: An extensible platform for privacy-preserving data analysis. In *SIGMOD Conference*. ACM, 19–30. DOI : <https://doi.org/10.1145/1559845.1559850>
- [180] Blaž Meden, Peter Rot, Philipp Terhörst, Naser Damer, Arjan Kuijper, Walter J. Scheirer, Arun Ross, Peter Peer, and Vitomir Štruc. 2021. Privacy-enhancing face biometrics: A comprehensive survey. *IEEE Trans. Inf. Forens. Secur.* 16 (2021), 4147–4183. DOI : <https://doi.org/10.1109/TIFS.2021.3096024>
- [181] Lubin Meng, Xue Jiang, Jian Huang, Wei Li, Hanbin Luo, and Dongrui Wu. 2023. User identity protection in EEG-based brain–computer interfaces. *IEEE Trans. Neural Syst. Rehab. Eng.* 31 (2023), 3576–3586. DOI : <https://doi.org/10.1109/TNSRE.2023.3310883>
- [182] Weizhi Meng, Duncan S. Wong, Steven Furnell, and Jianying Zhou. 2015. Surveying the development of biometric user authentication on mobile phones. *IEEE Commun. Surv. Tutor.* 17, 3 (2015), 1268–1293. DOI : <https://doi.org/10.1109/comst.2014.2386915>
- [183] Yan Meng, Yuxia Zhan, Jiachun Li, Suguo Du, Haojin Zhu, and Xuemin Shen. 2024. De-anonymizing avatars in virtual reality: Attacks and countermeasures. *IEEE Trans. Mob. Comput.* 23, 12 (2024), 13342–13357. DOI : <https://doi.org/10.1109/TMC.2024.3426046>
- [184] Sarina Meyer, Florian Lux, Pavel Denisov, Julia Koch, Pascal Tilli, and Ngoc Thang Vu. 2022. Speaker anonymization with phonetic intermediate representations. In *Interspeech Conference*. 4925–4929. DOI : <https://doi.org/10.21437/Interspeech.2022-10703>
- [185] Sarina Meyer, Florian Lux, Julia Koch, Pavel Denisov, Pascal Tilli, and Ngoc Thang Vu. 2023. Prosody is not identity: A speaker anonymization approach using prosody cloning. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 1–5. DOI : <https://doi.org/10.1109/ICASSP49357.2023.10096607>
- [186] Xiaoxiao Miao, Xin Wang, Erica Cooper, Junichi Yamagishi, and Natalia Tomashenko. 2022. Language-independent speaker anonymization approach using self-supervised pre-trained models. In *the Speaker and Language Recognition Workshop (Odyssey'22)*. 279–286. DOI : <https://doi.org/10.21437/Odyssey.2022-39>
- [187] Xiaoxiao Miao, Xin Wang, Erica Cooper, Junichi Yamagishi, and Natalia Tomashenko. 2023. Speaker anonymization using orthogonal householder neural network. *IEEE/ACM Trans. Audio, Speech Lang. Process* 31 (Sept. 2023), 3681–3695. DOI : <https://doi.org/10.1109/TASLP.2023.3313429>
- [188] Denis Migdal and Christophe Rosenberger. 2019. Keystroke dynamics anonymization system. In *International Joint Conference on e-Business and Telecommunications*. SCITEPRESS, 448–455. DOI : <https://doi.org/10.5220/0007923804480455>

- [189] Denis Migdal and Christophe Rosenberger. 2019. My behavior is my privacy & secure password! In *Conference on Cyberworlds*. IEEE, 299–307. DOI : <https://doi.org/10.1109/cw.2019.00056>
- [190] G. B. Moody and R. G. Mark. 1990. The MIT-BIH arrhythmia database on CD-ROM and software for use with it. In *Proceedings. Computers in Cardiology*. IEEE, 185–188. DOI : <https://doi.org/10.1109/cic.1990.144205>
- [191] Saemi Moon, Myeonghyeon Kim, Zhenyue Qin, Yang Liu, and Dongwoo Kim. 2023. Anonymization for skeleton action recognition. In *37th AAAI Conference on Artificial Intelligence and 35th Conference on Innovative Applications of Artificial Intelligence and 13th Symposium on Educational Advances in Artificial Intelligence (AAAI'23/IAAI'23/EAAI'23)*. AAAI Press, Article 1685, 9 pages. DOI : <https://doi.org/10.1609/aaai.v37i12.26754>
- [192] Asunción Moreno, Dolors Poch, Antonio Bonafonte, Eduardo Lleida, Joaquim Llisteri, José Mariño, and Climent Nadeu. 1993. Albayzin speech database: Design of the phonetic corpus. In *Eurospeech Conference*.
- [193] Aymen Mtibaa, Dijana Petrovska-Delacretaz, and Ahmed Ben Hamida. 2018. Cancelable speaker verification system based on binary Gaussian mixtures. In *Advanced Technologies for Signal and Image Processing*. IEEE, 1–6. DOI : <https://doi.org/10.1109/atsip.2018.8364513>
- [194] Naoya Mukojima, Masaki Yasugi, Yasuhiro Mizutani, Takeshi Yasui, and Hirotsugu Yamamoto. 2022. Deep-learning-assisted single-pixel imaging for gesture recognition in consideration of privacy. *IEICE Transactions on Electronics* E105.C, 2 (February 2022), 79–85. DOI : <https://doi.org/10.1587/transele.2021di0002>
- [195] Arsha Nagrani, Joon Son Chung, and Andrew Senior. 2017. VoxCeleb: A large-scale speaker identification dataset. In *Interspeech Conference*. CoRR abs/1706.08612
- [196] Vivek Nair, Wenbo Guo, James F. O'Brien, Louis Rosenberg, and Dawn Song. 2024. Deep motion masking for secure, usable, and scalable real-time anonymization of ecological virtual reality motion data. In *IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. 493–500. DOI : <https://doi.org/10.1109/VRW62533.2024.00096>
- [197] Vivek Nair, Wenbo Guo, Rui Wang, James F. O'Brien, Louis Rosenberg, and Dawn Song. 2024. Berkeley open extended reality recordings 2023 (BOXRR-23): 4.7 Million motion capture recordings from 105,000 XR users. *IEEE Transactions on Visualization and Computer Graphics* 30, 5 (May 2024), 2239–2246. DOI : <https://doi.org/10.1109/tvcg.2024.3372087>
- [198] Vivek Nair, Mark Roman Miller, Rui Wang, Brandon Huang, Christian Rack, Marc Erich Latoschik, and James F. O'Brien. 2024. Effect of data degradation on motion re-identification. In *IEEE 25th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*. 85–90. DOI : <https://doi.org/10.1109/WoWMoM60985.2024.00026>
- [199] Carol Neidle, Augustine Opoku, and Dimitris Metaxas. 2022. ASL Video Corpora and Sign Bank: Resources available through the American Sign Language Linguistic Research Project (ASLLRP). DOI : <https://doi.org/10.48550/ARXIV.2201.07899>
- [200] Alexandru Nelus and Rainer Martin. 2021. Privacy-preserving audio classification using variational information feature extraction. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 29 (2021), 2864–2877. DOI : <https://doi.org/10.1109/taslp.2021.3108063>
- [201] Alexandru Nelus and Rainer Martin. 2018. Gender discrimination versus speaker Identification through privacy-aware adversarial feature extraction. In *Speech Communication; 13th ITG-Symposium*, IEEE, 1–5.
- [202] Francesco Nespoli, Daniel Barreda, Jörg Bitzer, and Patrick A. Naylor. 2023. Two-stage voice anonymization for enhanced privacy. In *INTERSPEECH Conference*. 3854–3858. DOI : <https://doi.org/10.21437/Interspeech.2023-1341>
- [203] SUNY Downstate Medical Center Neurodynamics Laboratory. 1999. EEG Database. Retrieved from <http://kdd.ics.uci.edu/databases/eeg/eeg.data.html>
- [204] Alexis Nolin-Lapalme, Robert Avram, and Hussin Julie. 2023. PrivECG: Generating private ECG for end-to-end anonymization. In *Machine Learning for Healthcare Conference*. PMLR, 509–528.
- [205] Nymi. 2019. Always on authentication. Retrieved from <https://nyimi.com/>
- [206] Iyad Obeid and Joseph Picone. 2016. The Temple University Hospital EEG data corpus. *Front. Neurosci.* 10 (2016), 196.
- [207] Ikenna Odinaka, Po-Hsiang Lai, Alan D. Kaplan, Joseph A. O'Sullivan, Erik J. Sirevaag, and John W. Rohrbaugh. 2012. ECG biometric recognition: A comparative analysis. *IEEE Trans. Inf. Forens. Secur.* 7, 6 (Dec. 2012), 1812–1824. DOI : <https://doi.org/10.1109/tifs.2012.2215324>
- [208] Yoshitaka Ohshio, Haruka Adachi, Kenta Iwai, Takanobu Nishiura, and Yoichi Yamashita. 2018. Active speech obscuration with speaker-dependent human speech-like noise for speech privacy. In *Asia-Pacific Signal and Information Processing Association Annual Summit*. IEEE, 1252–1255. DOI : <https://doi.org/10.23919/apsipa.2018.8659754>
- [209] J. Ortega-Garcia, J. Fierrez-Aguilar, D. Simon, J. Gonzalez, M. Faundez-Zanuy, V. Espinosa, A. Satue, I. Hernaez, J.-J. Igarza, C. Vivaracho et al. 2003. MCYT baseline corpus: A bimodal biometric database. *Vis., Image Signal Process.* 150, 6 (2003), 395. DOI : <https://doi.org/10.1049/ip-vis:20031078>
- [210] Michele Panariello, Francesco Nespoli, Massimiliano Todisco, and Nicholas Evans. 2024. Speaker anonymization using neural audio codec language models. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4725–4729.

- [211] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. LibriSpeech: An ASR corpus based on public domain audio books. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 5206–5210. DOI : <https://doi.org/10.1109/icassp.2015.7178964>
- [212] Julien Pansiot, Danail Stoyanov, Douglas McIlwraith, Benny P. L. Lo, and G. Z. Yang. 2007. Ambient and wearable sensor fusion for activity recognition in healthcare monitoring systems. In *Workshop on Wearable and Implantable Body Sensor Networks*. Springer, 208–212. DOI : https://doi.org/10.1007/978-3-540-70994-7_36
- [213] Sree Hari Krishnan Parthasarathi, Herve Bourlard, and Daniel Gatica-Perez. 2011. LP residual features for robust, privacy-sensitive speaker diarization. In *Interspeech Conference*.
- [214] Sree Hari Krishnan Parthasarathi, H. Bourlard, and D. Gatica-Perez. 2013. Wordless sounds: Robust speaker diarization using privacy-preserving audio representations. *Trans. Audio, Speech, Lang. Process.* 21, 1 (Jan. 2013), 85–98. DOI : <https://doi.org/10.1109/tasl.2012.2215588>
- [215] Sree Hari Krishnan Parthasarathi, Mathew Magimai.-Doss, Daniel Gatica-Perez, and Hervé Bourlard. 2009. Speaker change detection with privacy-preserving audio cues. In *International Conference on Multimodal Interfaces*. ACM Press, 343. DOI : <https://doi.org/10.1145/1647314.1647385>
- [216] Damian Pascual, Alireza Amirshahi, Amir Aminifar, David Atienza, Philippe Ryvlin, and Roger Wattenhofer. 2021. EpilepsyGAN: Synthetic epileptic brain activities with privacy preservation. *IEEE Transactions on Biomedical Engineering* 68, 8 (August 2021), 2435–2446. DOI : <https://doi.org/10.1109/tbme.2020.3042574>
- [217] Manas A. Pathak and Bhiksha Raj. 2012. Privacy-preserving speaker verification as password matching. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 1849–1852. DOI : <https://doi.org/10.1109/icassp.2012.6288262>
- [218] Jose Patino, Natalia Tomashenko, Massimiliano Todisco, Andreas Nautsch, and Nicholas Evans. 2021. Speaker anonymisation using the McAdams coefficient. In *Interspeech Conference*. ISCA, 1099–1103. DOI : <https://doi.org/10.21437/Interspeech.2021-1070>
- [219] Juan M. Perero-Codosero, Fernando M. Espinoza-Cuadros, and Luis A. Hernández-Gómez. 2022. X-vector anonymization using autoencoders and adversarial training for preserving speech privacy. *Comput. Speech Lang.* 74, C (July 2022), 13 pages. DOI : <https://doi.org/10.1016/j.csl.2022.101351>
- [220] David I. Perrett, Sean N. Talamas, Patrick Cairns, and Audrey J. Henderson. 2020. Skin color cues to human health: Carotenoids, aerobic fitness, and body fat. *Frontiers in Psychology* 11 (March 2020). DOI : <https://doi.org/10.3389/fpsyg.2020.00392>
- [221] Ken Pfeuffer, Matthias J. Geiger, Sarah Prange, Lukas Mecke, Daniel Buschek, and Florian Alt. 2019. Behavioural biometrics in VR. In *CHI Conference on Human Factors in Computing Systems*. ACM, 1–12. DOI : <https://doi.org/10.1145/3290605.3300340>
- [222] Esteban Piacentino and Cecilio Angulo. 2020. Generating fake data using GANs for anonymizing healthcare data. In *International Work-Conference on Bioinformatics and Biomedical Engineering*. Springer, 406–417.
- [223] R. Plamondon and S. N. Srihari. 2000. Online and off-line handwriting recognition: A comprehensive survey. *IEEE Trans. Pattern Anal. Mach. Intell.* 22, 1 (2000), 63–84. DOI : <https://doi.org/10.1109/34.824821>
- [224] Kurt Plarre, Andrew Raji, Syed Monowar Hossain, Amin Ahsan Ali, Motohiro Nakajima, Mustafa Al’absi, Emre Ertin, Thomas Kamarck, Santosh Kumar, Marcia Scott et al. 2011. Continuous inference of psychological stress from sensory measurements collected in the natural environment. In *International Conference on Information Processing in Sensor Networks*. IEEE, ACM, 97–108.
- [225] M. Pobar and I. Ipsic. 2014. Online speaker de-identification using voice transformation. In *Convention on Information and Communication Technology, Electronics and Microelectronics*. IEEE, 1264–1267. DOI : <https://doi.org/10.1109/mipro.2014.6859761>
- [226] Bogdan Pogorelec, Zoran Bosnić, and Matjaž Gams. 2011. Automatic recognition of gait-related health problems in the elderly using machine learning. *Multim. Tools Applic.* 58, 2 (Nov. 2011), 333–354. DOI : <https://doi.org/10.1007/s11042-011-0786-1>
- [227] Frank E. Pollick, Jim W. Kay, Katrin Heim, and Rebecca Stringer. 2005. Gender recognition from point-light walkers. *J. Exp. Psychol. Hum. Percept. Perform.* 31, 6 (Dec. 2005), 1247–1265. DOI : <https://doi.org/10.1037/0096-1523.31.6.1247>
- [228] Alex Poole and Linden J. Ball. 2006. Eye tracking in HCI and usability research. In *Encyclopedia of Human Computer Interaction*. IGI Global, 211–219. DOI : <https://doi.org/10.4018/978-1-59140-562-7.ch034>
- [229] Jose Portelo, Alberto Abad, Bhiksha Raj, and Isabel Trancoso. 2013. Secure binary embeddings of front-end factor analysis for privacy preserving speaker verification. In *INTERSPEECH Conference*. 2494–2498.
- [230] Jose Portelo, Bhiksha Raj, Alberto Abad, and Isabel Trancoso. 2014. Privacy-preserving speaker verification using secure binary embeddings. In *Convention on Information and Communication Technology, Electronics and Microelectronics*. IEEE, 1268–1272. DOI : <https://doi.org/10.1109/mipro.2014.6859762>
- [231] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz et al. 2011. The Kaldi speech recognition toolkit. In *Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society.

- [232] Gauri P. Prajapati, Dipesh K. Singh, Preet P. Amin, and Hemant A. Patil. 2021. Voice privacy through x-vector and CycleGAN-based anonymization. In *Interspeech Conference*. ISCA, 1684–1688. DOI: <https://doi.org/10.21437/Interspeech.2021-1573>
- [233] Gauri P. Prajapati, Dipesh K. Singh, Preet P. Amin, and Hemant A. Patil. 2022. Voice privacy using CycleGAN and time-scale modification. *Comput. Speech Lang.* 74, C (July 2022), 30 pages. DOI: <https://doi.org/10.1016/j.csl.2022.101353>
- [234] Jiří Přibíl, Anna Přibilová, and Jindřich Matoušek. 2018. Evaluation of speaker de-identification based on voice gender and age conversion. *J. Electric. Eng.* 69, 2 (Mar. 2018), 138–147. DOI: <https://doi.org/10.2478/jee-2018-0017>
- [235] Jianwei Qian, Haohua Du, Jiahui Hou, Linlin Chen, Taeho Jung, and Xiangyang Li. 2021. Speech sanitizer: Speech content desensitization and voice anonymization. *IEEE Transactions on Dependable and Secure Computing* 18, 6 (November 2021), 2631–2642. DOI: <https://doi.org/10.1109/tdsc.2019.2960239>
- [236] Jianwei Qian, Haohua Du, Jiahui Hou, Linlin Chen, Taeho Jung, and Xiang-Yang Li. 2018. Hidebehind: Enjoy voice input with voiceprint unclonability and anonymity. In *Conference on Embedded Networked Sensor Systems*. ACM, 82–94. DOI: <https://doi.org/10.1145/3274783.3274855>
- [237] Jianwei Qian, Feng Han, Jiahui Hou, Chunhong Zhang, Yu Wang, and Xiang-Yang Li. 2018. Towards privacy-preserving speech data publishing. In *INFOCOM Conference*. IEEE, 1079–1087. DOI: <https://doi.org/10.1109/infocom.2018.8486250>
- [238] Yaron Rachlin and Dror Baron. 2008. The secrecy of compressed sensing measurements. In *Allerton Conference*. IEEE, 813–817. DOI: <https://doi.org/10.1109/allerton.2008.4797641>
- [239] Vibhor Rastogi and Suman Nath. 2010. Differentially private aggregation of distributed time-series with transformation and encryption. In *SIGMOD Conference*. ACM, 735–746. DOI: <https://doi.org/10.1145/1807167.1807247>
- [240] Vijay Ravi, Jinhan Wang, Jonathan Flint, and Abeer Alwan. 2024. Enhancing accuracy and privacy in speech-based depression detection through speaker disentanglement. *Comput. Speech Lang.* 86, C (June 2024), 24 pages. DOI: <https://doi.org/10.1016/j.csl.2023.101605>
- [241] S. R. S. Reddy, Sravani Nalluri, Subramanyam Kuniseti, S. Ashok, and B. Venkatesh. 2019. Content-based movie recommendation system using genre correlation. In *Smart Intelligent Computing and Applications*. Springer, 391–397.
- [242] Kenneth Revett, Hamid Jahankhani, Sérgio Tenreiro de Magalhães, and Henrique Santos. 2008. A survey of user authentication based on mouse dynamics. In *International Conference on Global e-Security*. Springer, 210–219.
- [243] Douglas A. Reynolds. 1995. Speaker identification and verification using Gaussian mixture speaker models. *Speech Commun.* 17, 1 (Aug. 1995), 91–108. DOI: [https://doi.org/10.1016/0167-6393\(95\)00009-d](https://doi.org/10.1016/0167-6393(95)00009-d)
- [244] Douglas A. Reynolds, Thomas F. Quatieri, and Robert B. Dunn. 2000. Speaker verification using adapted Gaussian mixture models. *Digit. Signal Process.* 10, 1 (Jan. 2000), 19–41. DOI: <https://doi.org/10.1006/dspr.1999.0361>
- [245] Slobodan Ribaric, Aladdin Ariyaeeinia, and Nikola Pavesic. 2016. De-identification for privacy protection in multimedia content: A survey. *Signal Process.: Image Commun.* 47 (Sept. 2016), 131–151. DOI: <https://doi.org/10.1016/j.image.2016.05.020>
- [246] Pierre Rougé, Ali Moukadem, Alain Dieterlen, Antoine Boutet, and Carole Frindel. 2022. Generalizable features for anonymizing motion signals based on the zeros of the short-time fourier transform. *J. Signal Process. Syst.* 95, 1 (July 2022), 89–99. DOI: <https://doi.org/10.1007/s11265-022-01798-9>
- [247] Zhang Rui and Zheng Yan. 2018. A survey on biometric authentication: Toward secure and privacy-preserving identification. *IEEE Access* 7 (2018), 5994–6009.
- [248] Napa Sae-Bae and Nasir Memon. 2013. A simple and effective method for online signature verification. In *2013 International Conference of the BIOSIG Special Interest Group (BIOSIG)*. 1–12.
- [249] Nazir Saleheen, Supriyo Chakraborty, Nasir Ali, Md Mahbubur Rahman, Syed Monowar Hossain, Rummana Bari, Eugene Buder, Mani Srivastava, and Santosh Kumar. 2016. MSieve: Differential behavioral privacy in time of mobile sensor data. In *International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. ACM, New York, USA, 706–717. DOI: <https://doi.org/10.1145/2971648.2971753>
- [250] Ben Saunders, Necati Cihan Camgoz, and Richard Bowden. 2021. Anonymsign: Novel human appearance synthesis for sign language video anonymisation. In *International Conference on Automatic Face and Gesture Recognition*. IEEE, 1–8. DOI: <https://doi.org/10.1109/FG52635.2021.9666984>
- [251] G. Schalk, D. J. McFarland, T. Hinterberger, N. Birbaumer, and J. R. Wolpaw. 2004. BCI2000: A general-purpose brain-computer interface (BCI) system. *IEEE Trans. Biomed. Eng.* 51, 6 (2004), 1034–1043. DOI: <https://doi.org/10.1109/TBME.2004.827072>
- [252] Abdur R. Shahid and Sajedul Talukder. 2021. Evaluating machine learning models for handwriting recognition-based systems under local differential privacy. In *Innovations in Intelligent Systems and Applications Conference*. IEEE, 1–6. DOI: <https://doi.org/10.1109/ASYU52992.2021.9598983>

- [253] Ali Shahin Shamsabadi, Brij Mohan Lal Srivastava, Aurélien Bellet, Nathalie Vauquier, Emmanuel Vincent, Mohamed Maouche, Marc Tommasi, and Nicolas Papernot. 2023. Differentially private speaker anonymization. *Proc. Privac. Enhanc. Technol.* 1 (2023), 98–114.
- [254] Dushyant Sharma, Francesco Nespola, Rong Gong, and Patrick A. Naylor. 2023. Canonical voice conversion and dual-channel processing for improved voice privacy of speech recognition data. In *31st European Signal Processing Conference (EUSIPCO)*. 66–70. DOI: <https://doi.org/10.23919/EUSIPCO58844.2023.10289777>
- [255] Michael Sherman, Gradeigh Clark, Yulong Yang, Shridatt Sugrim, Arttu Modig, Janne Lindqvist, Antti Oulasvirta, and Teemu Roos. 2014. User-generated free-form gestures for authentication. In *MobiSys Conference*. ACM, New York, NY, USA, 176–189. DOI: <https://doi.org/10.1145/2594368.2594375>
- [256] Md Shopon, Sanjida Nasreen Tumpa, Yajurv Bhatia, K. N. Pavan Kumar, and Marina L. Gavrilova. 2021. Biometric systems de-identification: Current advancements and future directions. *J. Cybersecur. Privac.* 1, 3 (2021), 470–495. DOI: <https://doi.org/10.3390/jcp1030024>
- [257] Dipesh K. Singh, Gauri P. Prajapati, and Hemant A. Patil. 2024. Voice privacy using time-scale and pitch modification. *SN Comput. Sci.* 5, 2 (Jan. 2024), 19 pages. DOI: <https://doi.org/10.1007/s42979-023-02549-8>
- [258] Girijesh Singh, Palak Patel, Muhammad Asaduzzaman, and Garima Bajwa. 2023. Selective EEG signal anonymization using multi-objective autoencoders. In *20th Annual International Conference on Privacy, Security and Trust (PST)*. 1–7. DOI: <https://doi.org/10.1109/PST58708.2023.10320167>
- [259] Vincent Sitzmann, Ana Serrano, Amy Pavel, Maneesh Agrawala, Diego Gutierrez, Belen Masia, and Gordon Wetstein. 2018. Saliency in VR: How do people explore virtual environments? *IEEE Trans. Vis. Comput. Graph.* 24, 4 (2018), 1633–1642. DOI: <https://doi.org/10.1109/TVCG.2018.2793599>
- [260] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur. 2018. X-vectors: Robust DNN embeddings for speaker recognition. In *Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 5329–5333. DOI: <https://doi.org/10.1109/icassp.2018.8461375>
- [261] Cristina Soaz and Klaus Diepold. 2016. Step detection and parameterization for gait assessment using a single waist-worn accelerometer. *Trans. Biomed. Eng.* 63, 5 (2016), 933–942. DOI: <https://doi.org/10.1109/TBME.2015.2480296>
- [262] Petr Sojka, Aleš Horák, Ivan Kopeček, and Karel Pala (Eds.). 2014. *Text, Speech and Dialogue Lecture Notes in Computer Science*, Vol. 8655. Springer International Publishing. DOI: <https://doi.org/10.1007/978-3-319-10816-2>
- [263] Lal Srivastava, Brij Mohan, Nathalie Vauquier, Md Sahidullah, Aurelien Bellet, Marc Tommasi, and Emmanuel Vincent. 2020. Evaluating voice conversion-based privacy protection against informed attackers. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, Barcelona, Spain, 2802–2806. DOI: <https://doi.org/10.1109/ICASSP40776.2020.9053868>
- [264] Ioanna-Ourania Stathopoulou and George A. Tsihrintzis. 2011. Emotion recognition from body movements and gestures. In *Intelligent Interactive Multimedia Systems and Services*. Springer, 295–303.
- [265] Julian Steil, Inken Hagedstedt, Michael Xuelin Huang, and Andreas Bulling. 2019. Privacy-aware eye tracking using differential privacy. In *International Symposium on Eye Tracking Research & Applications (ETRA)*. ACM, New York, NY, USA, 1–9. DOI: <https://doi.org/10.1145/3314111.3319915>
- [266] Nathan J. Stevenson, Karoliina Tapani, Leena Lauronen, and Sampsa Vanhatalo. 2019. A dataset of neonatal EEG recordings with seizure annotations. *Scient. Data* 6, 1 (2019), 1–8.
- [267] Tino Stöckel, Robert Jacksteit, Martin Behrens, Ralf Skripitz, Rainer Bader, and Anett Mau-Moeller. 2015. The mental representation of the human gait in young and older adults. *Front. Psychol.* 6 (2015), 943. DOI: <https://doi.org/10.3389/fpsyg.2015.00943>
- [268] Fahim Sufi, Seedahmed Mahmoud, and Ibrahim Khalil. 2008. A new ECG obfuscation method: A joint feature extraction & corruption approach. In *Conference on Information Technology and Applications in Biomedicine*. IEEE, 334–337. DOI: <https://doi.org/10.1109/itab.2008.4570644>
- [269] Shravani Sur and V. K. Sinha. 2009. Event-related potential: An overview. *Industr. Psychiat. J.* 18, 1 (2009), 70. DOI: <https://doi.org/10.4103/0972-6748.57865>
- [270] Cees H. Taal, Richard C. Hendriks, Richard Heusdens, and Jesper Jensen. 2010. A short-time objective intelligibility measure for time-frequency weighted noisy speech. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 4214–4217. DOI: <https://doi.org/10.1109/ICASSP.2010.5495701>
- [271] Takahiro Tamesue and Tetsuro Saeki. 2014. Sound masking for achieving speech privacy with parametric acoustic array speaker. In *Soft Computing and Intelligent Systems and Advanced Intelligent Systems*. IEEE, 1134–1137. DOI: <https://doi.org/10.1109/scis-isis.2014.7044805>
- [272] Pin Shen Teh, Andrew Beng Jin Teoh, and Shigang Yue. 2013. A survey of Keystroke dynamics biometrics. *The Scientific World Journal* 2013 (2013), 1–24. DOI: <https://doi.org/10.1155/2013/408280>
- [273] Pin Shen Teh, Ning Zhang, Andrew Beng Jin Teoh, and Ke Chen. 2016. A survey on touch dynamics authentication in mobile devices. *Comput. Secur.* 59 (June 2016), 210–235. DOI: <https://doi.org/10.1016/j.cose.2016.03.003>

- [274] Daksh Thapar, Chetan Arora, and Aditya Nigam. 2020. Is sharing of egocentric video giving away your biometric signature? In *Computer Vision – ECCV 2020*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer International Publishing, Cham, 399–416.
- [275] Daksh Thapar, Aditya Nigam, and Chetan Arora. 2021. Anonymizing egocentric videos. In *International Conference on Computer Vision (ICCV)*. IEEE, 2300–2309. DOI: <https://doi.org/10.1109/ICCV48922.2021.00232>
- [276] Ngoc-Dung T. Tieu, Huy H. Nguyen, Hoang-Quoc Nguyen-Son, Junichi Yamagishi, and Isao Echizen. 2017. An approach for gait anonymization using deep learning. In *Workshop on Information Forensics and Security*. IEEE, 1–6. DOI: <https://doi.org/10.1109/wifs.2017.8267657>
- [277] Ngoc-Dung T. Tieu, Huy H. Nguyen, Hoang-Quoc Nguyen-Son, Junichi Yamagishi, and Isao Echizen. 2019. Spatio-temporal generative adversarial network for gait anonymization. *J. Inf. Secur. Applic.* 46 (June 2019), 307–319. DOI: <https://doi.org/10.1016/j.jisa.2019.03.002>
- [278] Ngoc-Dung T. Tieu, Junichi Yamagishi, and Isao Echizen. 2020. Color transfer to anonymized gait images while maintaining anonymization. In *Asia-Pacific Signal and Information Processing Association Annual Symposium*. 1406–1413.
- [279] Natalia Tomashenko, Xin Wang, Xiaoxiao Miao, Hubert Nourtel, Pierre Champion, Massimiliano Todisco, Emmanuel Vincent, Nicholas Evans, Junichi Yamagishi, and Jean François Bonastre. 2022. *The VoicePrivacy 2022 Challenge Evaluation Plan*. arXiv preprint. Retrieved from <https://arxiv.org/abs/2203.12468>
- [280] Toda Tomoki, Ling-Hui Chen, Daisuke Saito, Fernando Villavicencio, Mirjam Wester, Zhizheng Wu, and Junichi Yamagishi. 2016. The Voice Conversion Challenge 2016 dataset. University of Edinburgh. School of Informatics. Centre for Speech Technology Research. DOI: <https://doi.org/10.7488/ds/1575>
- [281] Quang Nhat Tran, Benjamin P. Turnbull, and Jiankun Hu. 2021. Biometrics and privacy-preservation: How do they evolve? *Open J. Comput. Societ.* 2 (2021), 179–191. DOI: <https://doi.org/10.1109/ojcs.2021.3068385>
- [282] Nikolaus F. Troje. 2002. Decomposing biological motion: A framework for analysis and synthesis of human gait patterns. *J. Vis.* 2, 5 (Sept. 2002), 2. DOI: <https://doi.org/10.1167/2.5.2>
- [283] TypingDNA. 2019. Retrieved from <https://www.typingdna.com>
- [284] Anthony Ngozichukwuka Uwaechia and Dzati Athiar Ramli. 2021. A comprehensive survey on ECG signals as new biometric modality for human authentication: Recent advances and future challenges. *IEEE Access* 9 (2021), 97760–97802. DOI: <https://doi.org/10.1109/ACCESS.2021.3095248>
- [285] Tavish Vaidya and Micah Sherr. 2019. You talk too much: Limiting privacy exposure via voice input. In *Security and Privacy Workshops (SPW)*. IEEE, 84–91. DOI: <https://doi.org/10.1109/spw.2019.00026>
- [286] Michel Valstar, Jonathan Gratch, Björn Schuller, Fabien Ringeval, Denis Lalanne, Mercedes Torres Torres, Stefan Scherer, Giota Stratou, Roddy Cowie, and Maja Pantic. 2016. AVEC 2016: Depression, mood, and emotion recognition workshop and challenge. In *6th International Workshop on Audio/Visual Emotion Challenge (AVEC)*. Association for Computing Machinery, New York, NY, USA, 3–10. DOI: <https://doi.org/10.1145/2988257.2988258>
- [287] I. van der Linde, U. Rajashekar, A. C. Bovik, and L. K. Cormack. 2009. DOVES: A database of visual eye movements. *Spatial Vision*. 161–177 pages. Retrieved from <http://live.ece.utexas.edu/research/doves>
- [288] Gabriele Vassallo, Tim Van hamme, Davy Preuveneers, and Wouter Joossens. 2017. Privacy-preserving behavioral authentication on smartphones. In *International Workshop on Human-centered Sensing, Networking, and Systems*. ACM, 1–6. DOI: <https://doi.org/10.1145/3144730.3144731>
- [289] Voice Vault. 2019. VoiceVault Voice Biometric Authentication. Retrieved from <https://voicevault.com/>
- [290] Changsheng Wan, Li Wang, and Vir V. Phoha. 2019. A survey on gait recognition. *Comput. Surv.* 51, 5 (Jan. 2019), 1–35. DOI: <https://doi.org/10.1145/3230633>
- [291] Shuo Wang, Ming Jiang, Xavier Morin Duchesne, Elizabeth A. Laugeson, Daniel P. Kennedy, Ralph Adolphs, and Qi Zhao. 2015. Atypical visual saliency in autism spectrum disorder quantified through model-based eye tracking. *Neuron* 88, 3 (Nov. 2015), 604–616. DOI: <https://doi.org/10.1016/j.neuron.2015.09.042>
- [292] Tao Wang, Yushu Zhang, Shuren Qi, Ruoyu Zhao, Zhihua Xia, and Jian Weng. 2024. Security and privacy on generative data in AIGC: A survey. *ACM Comput. Surv.* 57, 4, Article 82 (Dec. 2024), 34 pages. DOI: <https://doi.org/10.1145/3703626>
- [293] Yijun Wang, Xiaogang Chen, Xiaorong Gao, and Shangkai Gao. 2017. A benchmark dataset for SSVEP-based brain–computer interfaces. *IEEE Trans. Neural Syst. Rehab. Eng.* 25, 10 (2017), 1746–1752. DOI: <https://doi.org/10.1109/TNSRE.2016.2627556>
- [294] Leon Willenborg and Ton de Waal. 2001. *Elements of Statistical Disclosure Control*. Springer New York. DOI: <https://doi.org/10.1007/978-1-4613-0121-9>
- [295] Ethan Wilson, Azim Ibragimov, Michael J. Proulx, Sai Deep Tetali, Kevin Butler, and Eakta Jain. 2024. Privacy-preserving gaze data streaming in immersive interactive virtual reality: Robustness and user experience. *IEEE Trans. Visualiz. Comput. Graph.* 30, 5 (2024), 2257–2268. DOI: <https://doi.org/10.1109/TVCG.2024.3372032>

- [296] Shun-Chi Wu, Peng-Tzu Chen, A. Lee Swindlehurst, and Pei-Lun Hung. 2019. Cancelable biometric recognition with ECGs: Subspace-based approaches. *IEEE Trans. Inf. Forens. Secur.* 14, 5 (May 2019), 1323–1336. DOI: <https://doi.org/10.1109/tifs.2018.2876838>
- [297] Danny Wyatt, Tanzeem Choudhury, and Jeff Bilmes. 2007. Conversation detection and speaker segmentation in privacy-sensitive situated speech data. In *INTERSPEECH Conference*.
- [298] Zhaoyang Xia, Yuxiao Chen, Qilong Zhangli, Matt Huenerfauth, Carol Neidle, and Dimitris Metaxas. 2022. Sign language video anonymization. In *Workshop on the Representation and Processing of Sign Languages*.
- [299] Shilin Xiao, Xiaoyu Ji, Chen Yan, Zhicong Zheng, and Wenyuan Xu. 2023. MicPro: Microphone-based voice privacy protection. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*. Association for Computing Machinery, New York, NY, USA, 1302–1316. DOI: <https://doi.org/10.1145/3576915.3616616>
- [300] Yanyu Xu, Yanbing Dong, Junru Wu, Zhengzhong Sun, Zhiru Shi, Jingyi Yu, and Shenghua Gao. 2018. Gaze prediction in dynamic 360° immersive videos. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. 5333–5342. DOI: <https://doi.org/10.1109/CVPR.2018.00559>
- [301] Sherif Yacoub, Steve Simske, Xiaofan Lin, and John Burns. 2003. Recognition of emotions in interactive voice response systems. In *EUROSPEECH Conference*.
- [302] Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald. 2019. CSTR VCTK Corpus: English Multi-speaker Corpus for CSTR Voice Cloning Toolkit (version 0.92). University of Edinburgh. The Centre for Speech Technology Research (CSTR). DOI: <https://doi.org/10.7488/ds/2645>
- [303] Roman V. Yampolskiy and Venu Govindaraju. 2010. Taxonomy of behavioural biometrics. In *Behavioral Biometrics for Human Identification*. IGI Global, 1–43. DOI: <https://doi.org/10.4018/978-1-60566-725-6.ch001>
- [304] Qing Yang, Tao Wang, Ning Su, Shifu Xiao, and Zoi Kapoula. 2012. Specific saccade deficits in patients with Alzheimer’s disease at mild to moderate stage and in patients with amnesic mild cognitive impairment. *J. Amer. Aging Assoc.* 35, 4 (May 2012), 1287–1298. DOI: <https://doi.org/10.1007/s11357-012-9420-z>
- [305] Yulong Yang, Gradeigh D. Clark, Janne Lindqvist, and Antti Oulasvirta. 2016. Free-form gesture authentication in the wild. In *Conference on Human Factors in Computing Systems*. ACM, 3722–3735.
- [306] Yang Yang, Yury Kartynnik, Yunpeng Li, Jiuqiang Tang, Xing Li, George Sung, and Matthias Grundmann. 2024. STREAMVC: Real-time low-latency voice conversion. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 11016–11020. DOI: <https://doi.org/10.1109/ICASSP48485.2024.10446863>
- [307] Jixun Yao, Qing Wang, Pengcheng Guo, Ziqian Ning, and Lei Xie. 2024. Distinctive and natural speaker anonymization via singular value transformation-assisted matrix. *IEEE/ACM Trans. Audio, Speech, Lang. Process.* 32 (2024), 2944–2956. DOI: <https://doi.org/10.1109/TASLP.2024.3407600>
- [308] Jixun Yao, Qing Wang, Yi Lei, Pengcheng Guo, Lei Xie, Namin Wang, and Jie Liu. 2023. Distinguishable speaker anonymization based on formant and fundamental frequency scaling. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 1–5. DOI: <https://doi.org/10.1109/ICASSP49357.2023.10095120>
- [309] Xin Yao and Senquan An. 2023. DP-VoicePub: Differential privacy-based voice publication. In *IEEE International Symposium on Circuits and Systems (ISCAS)*. 1–5. DOI: <https://doi.org/10.1109/ISCAS46773.2023.10182113>
- [310] Yue Yao, Josephine Plested, Tom Gedeon, Yuchi Liu, and Zhengjie Wang. 2019. Improved techniques for building EEG feature filters. In *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–6. DOI: <https://doi.org/10.1109/ijcnn.2019.8852302>
- [311] Xin Xu Shaoji Zhang Ming Li Yao Shi, and Hui Bu. 2015. AISHELL-3: A multi-speaker Mandarin TTS corpus and the baselines. Retrieved from <https://arxiv.org/abs/2010.11567>
- [312] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven C. H. Hoi. 2021. Deep learning for person re-identification: A survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 6 (June 2022), 2872–2893. DOI: <https://doi.org/10.1109/tpami.2021.3054775>
- [313] Dit-Yan Yeung, Hong Chang, Yimin Xiong, Susan George, Ramanujan Kashi, Takashi Matsumoto, and Gerhard Rigoll. 2004. SVC2004: First international signature verification competition. In *Biometric Authentication*, David Zhang and Anil K. Jain (Eds.) (*Lecture Notes in Computer Science*, Vol. 3072). Springer Berlin, 16–22. DOI: https://doi.org/10.1007/978-3-540-25948-0_3
- [314] In-Chul Yoo, Keonnyeong Lee, Seonggyun Leem, Hyunwoo Oh, Bonggu Ko, and Dongsuk Yook. 2020. Speaker anonymization for personal information protection using voice conversion techniques. *IEEE Access* 8 (2020), 198637–198645. DOI: <https://doi.org/10.1109/ACCESS.2020.3035416>
- [315] Galit Yovel and Alice J. O’Toole. 2016. Recognizing people in motion. *Trends Cognit. Sci.* 20, 5 (May 2016), 383–395. DOI: <https://doi.org/10.1016/j.tics.2016.02.005>
- [316] Ruibin Yuan, Yuxuan Wu, Jacob Li, and Jaxter Kim. 2022. DeID-VC: Speaker de-identification via zero-shot pseudo voice conversion. In *Interspeech Conference*. 2593–2597. DOI: <https://doi.org/10.21437/Interspeech.2022-11036>
- [317] Emna Kalai Zaghouni, Adel Benzina, and Rabah Attia. 2017. ECG based authentication for e-healthcare systems: Towards a secured ECG features transmission. In *Wireless Communications and Mobile Computing Conference (IWCMC)*. IEEE, 1777–1783. DOI: <https://doi.org/10.1109/iwcmc.2017.7986553>

- [318] Emna Kalai Zaghouni, Adel Benzina, and Rabah Attia. 2017. ECG biometric template protection based on secure sketch scheme. In *Conference on Software, Telecommunications and Computer Networks (SoftCOM)*. IEEE, 1–5. DOI: <https://doi.org/10.23919/softcom.2017.8115526>
- [319] Mohammad-Reza Zare-Mirakabad, Fatemeh Kaveh-Yazdy, and Mohammad Tahmasebi. 2013. Privacy preservation by k-anonymizing ngrams of time. In *ISC Conference on Information Security and Cryptology*. 1–6. DOI: <https://doi.org/10.1109/ISCISC.2013.6767335>
- [320] Gao Zhang, Zhiwei Guan, Guozhong Dai, and Xiangshi Ren. 1998. A comparison of four interaction modes for CAD systems. In *Asia-Pacific Conference on Computer Human Interaction (APCHI)*. 82–87. DOI: <https://doi.org/10.1109/APCHI.1998.704160>
- [321] Guanglin Zhang, Sifan Ni, and Ping Zhao. 2020. Enhancing privacy preservation in speech data publishing. *Internet Things J.* 7, 8 (Aug. 2020), 7357–7367. DOI: <https://doi.org/10.1109/jiot.2020.2983228>
- [322] Ni Zhang and Yoshinori Yaginuma. 2012. A privacy-preserving and language-independent speaking detecting and speaker diarization approach for spontaneous conversation using microphones. In *International Conference on Signal Processing*. IEEE, 499–502. DOI: <https://doi.org/10.1109/icosp.2012.6491534>
- [323] Jianwei Zheng, Jianming Zhang, Sidy Danioko, Hai Yao, Hangyuan Guo, and Cyril Rakovski. 2020. A 12-lead electrocardiogram database for arrhythmia research covering more than 10,000 patients. *Scient. Data* 7, 1 (Feb. 2020). DOI: <https://doi.org/10.1038/s41597-020-0386-x>
- [324] Nan Zheng, Aaron Paloski, and Haining Wang. 2016. An efficient user verification system using angle-based mouse movement biometrics. *IEEE Trans. Inf. Forens. Secur.* 18, 3 (Apr. 2016), 1–27. DOI: <https://doi.org/10.1145/2893185>
- [325] Shuai Zheng, Junge Zhang, Kaiqi Huang, Ran He, and Tieniu Tan. 2011. Robust view transformation model for gait recognition. In *International Conference on Image Processing*. IEEE, 2073–2076. DOI: <https://doi.org/10.1109/icip.2011.6115889>
- [326] Wei-Long Zheng and Bao-Liang Lu. 2015. Investigating critical frequency bands and channels for EEG-based emotion recognition with deep neural networks. *Trans. Auton. Ment. Dev.* 7, 3 (2015), 162–175. DOI: <https://doi.org/10.1109/TAMD.2015.2431497>
- [327] Yu Zhong and Yunbin Deng. 2015. A survey on keystroke dynamics biometrics: Approaches, advances, and evaluations. In *Gate to Computer Science and Research*. Number 1. Science Gate Publishing P.C., 1–22. DOI: <https://doi.org/10.15579/gcsr.vol2.ch1>
- [328] Mohammad Zohaib. 2018. Dynamic difficulty adjustment (DDA) in computer games: A review. *Advances in Human-Computer Interaction* 2018 (November 2018), 1–12. DOI: <https://doi.org/10.1155/2018/5681652>

Received 11 September 2023; revised 19 March 2025; accepted 2 April 2025